

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 June 2002 (13.06.2002)

PCT

(10) International Publication Number
WO 02/46914 A2

(51) International Patent Classification⁷: G06F 9/00

(21) International Application Number: PCT/IL01/01133

(22) International Filing Date: 6 December 2001 (06.12.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/251,918 7 December 2000 (07.12.2000) US
PCT/IL01/00208 5 March 2001 (05.03.2001) IL

(71) Applicant (for all designated States except US): ADUVA INC. [US/US]; 19731 La Mar Drive, Cupertino, CA 95014 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): TE'ENI, Moddi

[IL/IL]; 3 Wiezel Street, 64241 Tel Aviv (IL). SHUFER, Ilan [IL/IL]; 12 Hachita Street, 50200 Bet Dagan (IL).

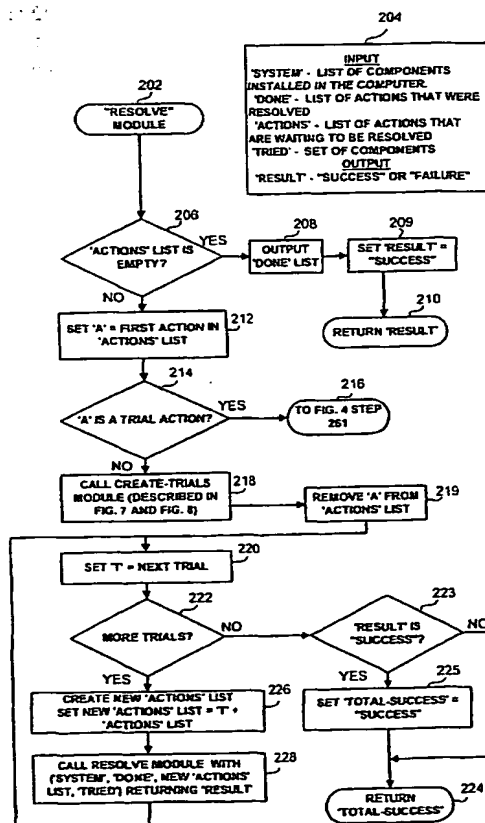
(74) Agents: AGMON, Jonathan et al.; Soroker - Agmon, 12th floor, Levinstein Tower, 23 Petach Tikva Road, 66184 Tel Aviv (IL).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: IMPROVED METHOD FOR RESOLVING DEPENDENCY CONFLICTS AMONG MULTIPLE OPERATIVE ENTITIES WITHIN A COMPUTING ENVIRONMENT



(57) Abstract: An improved method to set up automatically a computer system configuration in association with a system upgrade that involves the installation of a set of operative elements necessary for the reliable operation of the same. The correct, operative combination of hardware, system and application software components is accomplished by inter-component dependency checking and improved conflict resolving procedures in association with specific dependency rules tables and an ordered database that stores definitions of components and relationships thereof.

WO 02/46914 A2



European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *without international search report and to be republished upon receipt of that report*

IMPROVED METHOD FOR RESOLVING DEPENDENCY CONFLICTS AMONG MULTIPLE OPERATIVE ENTITIES WITHIN A COMPUTING ENVIRONMENT

5

RELATED APPLICATIONS

This application is a continuation in part of application U.S. Ser. No. 09/585,694 filed June 1st 2000. This application is related to application U.S. Ser. No. 09/586,685 filed June 1st 2000, and to application U.S. Ser. No 09/585,685 filed June 1st 2000.

10

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

The present invention relates generally to a method that establishes a correct combination of different hardware and software entities necessary for the reliable operation of a computer system and more particularly to an improved method for resolving dependency conflicts across diverse sets of functional entities while installing or removing specific operative elements in a computing environment.

DISCUSSION OF RELATED ART

20

Computing environments comprises of hardware components and software components. The components are divided into various types: processor devices, disk devices, printer devices, software packages, software libraries, kernel bases, kernel parameters, kernel modules, drivers, configuration files, flags, application packages and the like. The software elements also referred to as modules, or programs, comprise a multitude of executable instructions in hardware-readable format. The components are having diverse functionality or applications. Hardware devices are analog interfaces to the real world and perform physical work such as switching electronic circuits, transmitting electrical signals, magnetizing coils, pressing ink into paper surfaces, and the like while software modules operate in a digital mode and control both the operation of the

30

hardware devices, the operation of other software modules in a preset manner, supply various services involving data manipulation and computation by receiving, processing, transforming and outputting information units in a predefined manner. System software modules such as kernel modules supervise,
5 manage and control the operations of the computing environment in its entirety.

Typically in a computing environment various hardware and software entities operate in close co-operation. Therefore, within a typical computing environment a plurality of hardware and software entities is having multiple dependency relationships with a plurality of other hardware and software entities.
10 Such a dependency relationship could be defined as follows: in an exemplary computing environment for component 'A' to perform correctly, component 'B' and component 'C' are both needed. Therefore, in the exemplary computing environment such as an operating system platform the utilization of component 'A' necessitates the presence of component 'B' and the presence of component
15 'C'.

Computer systems are practically never static. Frequently changes have to be made. Changing requirements of the environment necessitates the addition and/or replacement of hardware devices or software modules thereby inducing the addition and/or replacement of other software modules. Improved versions of
20 present devices or modules effect the replacement of older versions thereof. New, state-of-art software packages are installed repeatedly in a dynamic process of development and growth. Modifying a computer system typically involves installation or removal of hardware devices or software modules - a process referred to as upgrading the system. When performing the predefined procedures
25 necessary for an upgrade to be implemented frequently dependency conflicts arise among the components present and the components to be installed. Such dependency conflicts turn the upgrading process into a complicated, prolonged, difficult, and sometimes unsuccessful operation.

Conventionally, users utilizing specific installation utilities perform the
30 installation and the update of components. The majority of utilities operate in a

basic fashion; installable components are installed, replacing, if necessary, present components, no dependency checks are made therefore, no attempts are made to solve the dependency conflicts.

5 Some more advanced utilities perform dependency checks, typically abort the installation process when required and also inform the users in regard to the conflicts but make no attempts to solve the related problem.

Formerly the present applicant submitted several patent applications including a method and system operative in resolving and preventing conflicts occurring during software installation within a computing system. The above-mentioned applications are listed hereinabove in the Related Applications section
10 of the present application. The above-mentioned applications suggest a system and method for the substantial resolution and prevention of potential operational conflicts between diverse software components being installed. The present invention is substantially based on the method and system offered by the above-mentioned prior applications. The present invention proposes the utilization of
15 several sub-methods to provide for enhanced execution times in certain computing environments wherein the timing factor is substantially important or even critical. The sub-methods are introduced in the form of additional executable software modules designed to be installed in the computing environment such that
20 the activation of the extra modules is achieved by appropriate modifications in the main logic module of the method. For example, specific "call" instructions that contain the memory addresses of the sub-modules thereby made operative in the loading, the initialization and the activation of the additional sub-methods could be introduced at suitable address locations in the instruction sequence of the main
25 logic module.

It will be obvious to those skilled in the art that there is a long felt need for a comprehensive, totally automated installation utility to assist users of small and medium-size computer system platforms in the exacting task of managing complex and dynamically evolving computing systems. Specifically there is an
30 urgent need for an effective installation utility designed to resolve automatically

inter-component dependency conflicts so as to enable routinely the accomplishment of fast, solid, efficient and convenient computer system upgrades.

5 SUMMARY OF THE PRESENT INVENTION

One aspect of the present invention regards a method of supporting users of computer client systems by providing assistance in the performance of a computer system upgrade procedure by automatically establishing an operatively correct combination of a set of components in a computing environment. The method consists of obtaining control tables from a storage device, creating result
10 tables on the storage device, examining control tables in order to identify potential dependency conflicts arising among said components, resolving said dependency conflicts, and notifying said user in regard to the content of said result tables in order to enable said user to perform a real system upgrade process.

15 The method of the present invention further consists of obtaining a system information table comprising a set of operative components, obtaining a first system upgrade, obtaining a component object, obtaining a component database, obtaining a rule table comprising installation related dependencies rules. The method further consists.

20 Further aspects of the present invention will be readily understood from the description which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully
25 from the following detailed description taken in conjunction with the drawings in which:

Fig. 1 is a block diagram of the computing environment in which the improved dependency conflict resolver module is operating, in accordance with a preferred embodiment of the present invention; and

Fig. 2 is a flow chart illustrating the flow of program logic of the resolve module, in accordance with the preferred embodiment of the present invention; and

5 Fig. 3 is a flowchart illustrating the flow of program logic of the "relative-system-contains-X" module, in accordance with the preferred embodiment of the present invention; and

Fig. 4 and Fig. 5 are flow charts illustrating the operation of the program module that handles the non-trial actions, in accordance with the preferred embodiment of the present invention; and

10 Fig. 6 is a flow chart illustrating the flow of program logic of the "checking-of-xor-rules" module, in accordance with the preferred embodiment of the present invention; and

Fig. 7 is a flow chart illustrating the flow of program logic of the "create-trials-for-an-install-action" module, in accordance with the preferred embodiment of the present invention; and

15 Fig. 8 is a flow chart illustrating the flow of program logic of the "create-trials-for-an-update-action" module, in accordance with the preferred embodiment of the present invention; and

Fig. 9 is a flow chart illustrating the flow of program logic of the "check-add-remove-rules" module, in accordance with the preferred embodiment of the present invention; and

20 Fig. 10 is a flow chart illustrating the flow of program logic of the "check-step" module, in accordance with the preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention overcomes the disadvantages of the prior art by providing an improved method and a system to improved dependency relationship analysis and to improved dependency conflict resolving among multiple sets of hardware and software components in a computing environment.

The present invention provides an improved conflict resolving mechanism implemented as a software program that associates lists of installed components with lists of components to be installed, removed or updated. By utilizing a set of related predetermined dependency rules pertaining to both components lists the dependency conflict resolving module resolves existing conflicts among the hardware and software components and guides the user through a sequence of appropriate actions. Consequently, the user can utilize the information so obtained in a useful manner such as performing one or more recommended installation procedures in order to accomplish a successful system upgrade.

It should be clearly understood that the description of the computing environment as a totality, the associated tables, programs and methods are provided in the text of this document for the purpose of enabling a better understanding of the present invention, and not intended as any limitation on the potential uses that will utilize the fundamental concepts of the disclosure within the scope of the attached claims.

Referring now to the drawings, in which like numerals represent elements throughout the several figures, aspects of the present invention and the exemplary operating environment will be described for purpose of clear understanding.

Fig. 1 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the present invention may be implemented. While the invention will be described in the general context of a client-based computer system, those skilled in the art will

recognize that the invention also may be implemented in combination with other environments.

Referring now to Fig. 1 there is shown constructed and operative in accordance with the preferred embodiment of the present invention the computing environment in which the improved Dependency Conflict Resolver Module is
5 operating.

System 10 is a computing platform controlled by an operating system 18. System 10 contains a storage device 12 such as a hard disk, a Central Processing Unit (CPU) 14, an I/O device 16, a Graphical User Interface (GUI) 20
10 such as display terminal, and a communication device 17 such as a modem that links system 10 via appropriate communication lines to a data network 11. There are two distinct sets of software entities involved in the process of dependency conflict resolving (1) control tables 31, and (2) application programs 21. Additionally system 10 contains a component objects table 44.

15 The component objects table 44 contains the collection of components to be installed. The components could be of all the practicable types typically comprising hardware and software elements of a computing environment; hardware devices, hardware drivers, operating system components such as kernels, utilities, programming languages, application programs, debuggers,
20 communication packages, application packages, libraries and the like. In the preferred embodiment of the present invention Component objects 44 is downloaded by communication device 17 from a central support server system and stored on storage device 12 such as a hard disk or the like. In different embodiments, other methods could be used For example component objects 44
25 may be downloaded automatically from the supporting server system after Conflict Resolver Module 40 completed the dependency resolving process.

The set of the control tables comprises original system information table 22, system upgrade table 24, and Knowledge Base 26 combining Component Data database 28, Rules tables 30, and Cost table 37. Rules tables 30 contains Xor

Rules table 32, Add-Remove Rules table 34, Handler Rules table 35, and Working Rules table 33.

Original system information table 22 contains the entire list of hardware and software components present in the operative system and stored on storage device 12. The user of system 10, utilizing user interface (GUI) 20, creates original system information table 22 by using specific predefined utility programs that access specific predefined system libraries and system files.

System Upgrade Table 24 is created and maintained by the user via user interface (GUI) 20 using specific programs. System Upgrade table 24, stored on storage device 12, is intended to hold the list of components to be installed.

Knowledge Base 28 is a set of related tables such as Component Data table 28, Xor Rules table 32 and Add-Remove Table 34, Handler rules table 35, and Working rules table 33.. Component Data 28 comprises useful data fields as well as control information concerning the components extant in component objects table 44. Additionally component data table 28 holds the requisite data pointers to the relevant entries in rules tables 30, specifically in Xor rules table 32 and in Add-Remove rules 34. Xor rules table 32 and Add-remove rules table 34 contain the dependency rules, encoded into a particular format by using a specific rules language, that refer to the component objects to be installed. The Knowledge base 26 is downloaded via communication device 17 from the support server system. Typically, only the requisite parts of Knowledge base 26 that are relevant to the current installation process or the current operative system are downloaded. The functionality of the Handler rules table 35, the Working rules table 33, and the Cost table 37 will be described hereinafter in association with the proposed sub-modules operative in the optimization of the execution time.

A Virtual Upgrade process is defined as the detailed and careful examination of the component data and the dependency rules associated with the component objects needed to be inserted into the existing framework of the operative system installed for potential dependency conflicts. In addition, where a dependency conflict is recognized the virtual upgrade process will effect the

resolving of practically all the potential dependency conflicts. In contrast, a Genuine Upgrade process is defined as the actual installation of the required components in the operative system after all the dependency conflicts had been resolved.

5 The dependency rules are inherited along a downward path of the Knowledge Base 26. Each component inherits the rules of the components above. For example, if a specific software utility such as an editor has a "need-rule" such as "the editor needs graphical user interface" then all the versions of same editor stored in the nodes under the editor in the knowledge base 26 will be having the
10 same "need-rule" implicitly, like as "editor version 3.1.1 needs graphical user interface" without having to be defined explicitly.

 In this document the term "upgrade" refers both to the installation and the removal of operative system components such as hardware devices, software modules, hardware drivers, application packages, operating system modules,
15 kernel bases, libraries and the like.

 The set of application programs or application modules comprises Virtual Upgrade Module 36, Conflict Resolver Module 40 and Component Object
Virtual Upgrade Module 36 is the manager of the virtual
20 upgrade process activated by the user via user interface (GUI) 20. Virtual Upgrade module 36 creates upgrade processes for the following tasks that are performed sequentially: (a) to collect all the information necessary for the dependency analysis and the dependency conflicts resolving process such as the relevant component data information units from component data table 28, the encoded dependency rules from xor-rules table 32 and from add-remove table 34
25 module, (b) to activate Conflict Resolver module 40 in order to check for potential dependency conflicts and to resolve said dependency conflicts that might arise as a result of the planned installation process, (c) to initiate Downloader Manager module in order to download the required component objects from the supporting server system or from any other source, and (d) to perform a genuine
30 installation of the selected component objects from component object table 44.

The intention of above general description was to provide a clear picture of the standing and functions of the Conflict Resolver Module 40 in the larger framework of the computing environment thereof. It will be easily perceived that Conflict Resolver Module 40 is the central element of the virtual
5 upgrade process. The operation and the functionalities of the conflict-resolver module 40 will be described next in association with the related control tables when necessary.

The elements, tables, operations, and rules set forth in the following description are described in full detail in the above-mentioned related patent
10 applications U.S. Ser. No. 09/586,685 filed June 1st 2000, and U.S. Ser. No. 09/585,694 filed June 1st 2000.

The input data elements for the conflict-resolver module will be described next. To prepare for the installation of the new components or for the upgrading of the operatively installed components an Original System Information
15 Table 22 is be created. In the preferred embodiment of the present invention, the user of the computing environment creates the original system table 22. It would be easily perceived that the creation of original table 22 could be initiated by other means as well such as automatic activation by appropriate software modules and the like. Original system 22 is built by utilizing diverse software application
20 programs that that extract the required information from the operatively installed system and insert the information into a tree-like data structure. The components installed in the operative system will be represented in the nodes of the tree-like structure designated by predefined numbers. The numbers will be held in the leaves of the tree-like data structure. Therefore, the original system information
25 table 22 comprises structured information in predefined format concerning the set of the components installed in the computer prior to a preferred installation process.

Installable components are either new components or new versions of existing components that are stored in the Component Objects Table 44 either in
30 object code form or in a source language. The Component Data table 28 that holds

a list of all the installable components points to the components stored in the Component Objects table 44. Component Data Table 28 is built as a tree-like data structure in a similar manner to Original system information table 22 and the information concerning the installable components is represented by predefined numbers therein. The numbers are held in the leaves of the Component Data Table 28 in a similar manner to the information stored in Original system information table 22. The numbering system is the same for both sets of information. Therefore the information in the leaves of the Component Data Table 28 makes available the recognition whether a specific component represented in the component data table 28 exists in the operatively installed system through accessing the original system information table 22. Apart from the information about installable components the nodes of the Component Data table 28 could hold other type of information such as various generic definitions referring to the installable components placed under them and pointed to by them. For example, a node could hold generic information such as "version", meaning that all leaves pointed to by the generic node are holding information about different versions of the same component.

The user can browse the Component Data Table 28 and user preferences can be set on any of the nodes therein. For example, the preferences set option makes available to the user the locking of a specific node to prevent the replacement thereof. It should be easily perceived that the replacement of a specific node (upgrade or downgrade i.e., replacing the component pointed to by the node with a higher or lower version of the same) could be caused not only by a particular upgrade thereof by also by the replacement of some other node and the consequent conflict-resolving series of actions of the resolver module discussed. The default preference is no update meaning that every component pointed to by this node will not be updated unless the user changes the preference.

The System Upgrade Information Table 24 is built from diverse sources. The table 24 contains (a) all the components that were selected by the user to be upgraded or installed (using the Component Data Table 28 either as a

default or as a preference) (b) components that were not installed or upgraded previously (in reference to the Original system information table 22) and (c) entirely new components that will be installed unless the user locked the nodes thereof in the component data table 28. As a result, the system upgrade table 24
5 comprises components that should be installed or upgraded in the operative system.

The system upgrade table 24 consists of a list of actions to be done. The list of actions built in the following form: "install 1024" or "update 1057 to 1058". An action is a request for installing and/or uninstalling. The list of actions
10 is defined as single actions to be resolved by the Conflict Resolver Module 40.

A single action on the list of actions comprises "action-type" or "what is to be done" and "component type" or "to whom to do what is to be done". A component-type can be a list such as "install 1088, 1089, 1123". To apply the action type on the component type a component type-specific rules-based
15 dependency check has to be made. Each component type has a set of rules that must be satisfied in order for the component type to be installed. The rules are required regarding the consistence of the system after the preferred installation. The rules are based on sets of components or predefined relationships between of components. For example, "Mail 4.3.2" depends on "mail daemon". Rules have
20 different "importance factor definition" or "importance weight" such as a "MUST" in which case the component will not be installed unless the rule is satisfied or a "RECOMMENDED" in which case the rule is not enforced automatically but left to the user's discretion. The rules are held in the Rules Tables 30. The xor rules are stored in XOR Rules 32 and the add-remove rules are
25 stored in the Add-Remove Rules 34. Both types of rules were described previously.

The sets of dependencies are lists of components with the attribute of either "XOR" (exactly one of) or "OR" (at least one of). The list of components can be recursive, i.e.; it could be a list of lists. For example, Editor = OR (vi, vim,
30 emacs), Vim = XOR (vim4, vim5), Vim5 = XOR (vim5.1, vim5.2, vim5.3) and

the like. The sets have inner priorities in order to enable choosing one of the components if the rules do not specify a specific one. Typically the priorities will be determined by the order of the different version that is a component with a higher version number will have priority over the same component with a lower version number.

Following the execution of the dependency checks the inter-component conflicts that were recognized has to be resolved. The resolving of the conflicts is accomplished by extending the action list. Each action has relevant rules, rules that might affect the possibility of achieving the action. For example, action "install component Z" might have a relevant rule such as "component Z needs component B" or "xor (component Z, component A)". In the first example installation of the component Z requires the installation of the component B as well. In the second example, the installation of the component Z requires removal of the component A.

If the dependency checks result indicate that component Z, which has to be installed, needs for the correct operation thereof component A, then the conflict could be resolved only by installing component A as well. Accordingly, the action list would be extended with the action type/component type "install A" and the appropriate dependency checks would be executed on the newly added action by using the component A-specific rules. In such a manner, the action list is extended until no more conflicts are left to be resolved.

Therefore, the dependency check uses a recursively activated process. Each step takes the set of operative components, the related operative actions that resulted in said set of operative components, and a new action that should be performed. After the performance of the current step, the output from the step could be an extended set of operative components, an extended operative action list and a flag indicating "success" or "failure" of the prior step. A log file is managed for each step to display to the user the various stages leading up to the installation and/or upgrade of the originally requested components and the

additional components installed and/or upgraded as necessary for conflict resolving.

The conflict-resolving module will be described next. Four lists are utilized as input data for the module. (1) The list of components in the original system information table 22 or the components installed before the preferred installation. (2) The set of actions to be done where an action is a request for installing, uninstalling or updating. (3) Rules regarding the consistence of the preferred system from the Rules Tables 30. (4) Sets of components that can apparently replace each other from the Component Data Table 28.

In the case of the successful completion of the conflict-resolving module, the output comprises a list and an indicator. (1) A list of all the possibilities to perform the task of preferred update. Each item on the list is a sequence of actions to be performed one after the other. (2) An indicator that indicates success. In case of an unsuccessful completion, the module returns an indicator indicating failure.

Each action on the action list could be a complex action like having more than one component type when using the action type "install". For example, in the action "install component A, component B, component C uninstall component D, component E, component F" the action will be divided into three parts defined as "trials" such as "install A uninstall D, E, F", "install B uninstall D, E, F" and the like. The trials are added to the action list replacing the original complex action with a specific flag added to the action record to indicate trial action.

For each action, the rules relevant to the action type and component type are checked. The rules checking process will be explained by using the example of the action "install X, uninstall Y, Z". First, the xor rules are checked. The xor set of component X has to be found and extracted from the Xor Rules 32. If the xor set of component X exists then the set comprises a set of components that only one of them could be operative in a consistent system. Consequently, if component X required to be installed all the other components in the xor set

thereof must be removed. The implementation of the xor rule check is to alter the action list by adding the components of the xor set to the "uninstall" part. Therefore, all the components from the xor set of X, except X itself, that are on the list of components, i.e., installed in the operative system and not already in the
5 uninstall part of the action list, are added to the said part of the said list.

The add-rules checking is performed by finding all the rules from Add-Remove Rules 34 concerning component X such as "component X needs component A or component B or not component N or not component M". If none of the necessary components in the rule such as A or B are in the list of the
10 operative system components and said list contains N and M, an action is created in the form of "install A or B or...". The list of actions thus created is called the "pre-action list", as all these actions should be performed before installing X.

For each of the components to uninstall remove checking is performed. Rules regarding uninstalling are extracted from Add-Remove Rules 34. These
15 rules are of the kind "component A needs component Y or component W or not component M", "component B needs component Z" and the like. For each such rule the preferred system should stay in a consistent state after the performance of respective install and remove-rule-resulting uninstall actions. To describe it more in detail:

20 "A is installed in the operative system, or will be installed by the action resulting from 'X needs A'."

b) "All the components in the needed part of the remove-rule (Y, W) are not installed in the operative system or are in the uninstall part of the action list."

25 c) "The 'or not' part components are installed in the operative system."

To satisfy the rule after the performance of the action a "post-action list" is created. The list comprises actions that should be performed after the action list. Therefore, and "update A" action is inserted into the post-action list.

In the following description of the program logic and the sequentially
30 or recursively executed sets of operative steps necessary to resolve practically all

the component-related dependency conflicts, specific data structures used as storage means for specific computer-related abstractions such as constants, variables, lists and the like. For the clarity of the disclosure, designations of said data structures appear in the text of the description in upper case characters. The designations of the same data structures will appear on the associated drawing enclosed by single quotes. For example, the dynamic data structure holding the list of actions to be executed in order to perform a requested system upgrade will be specified as ACTIONS in the text of the description and as 'ACTIONS' on the related drawings. Occasionally in order to provide for a more extended and clearer description, ACTIONS will be referred to as ACTIONS list, or a list of ACTIONS, and the like. The specific values given to data structures will appear both in the text of the description and on the drawings enclosed by double quotes such as "success", "failure", and the like.

Referring now to Fig. 2 illustrating the flow of program logic of the Resolve Module 202. The input and output data elements of the module 202 are held in allocated physical locations of the storage device 12 and are listed in frame 204. The input consists of the following elements: (1) SYSTEM which is a list of components installed in the computer, (2) DONE which is the list of actions that were resolved, (3) ACTIONS which is the list of actions that are waiting to be resolved, and (4) TRIED which is set of components. Initially, the DONE list is empty and ACTIONS is a list holding all the actions to be resolved. During the repeated iterations or recursive calls of module 202, DONE is being filled up with the resolved actions previously held in ACTIONS and copied from ACTIONS after the resolving thereof. TRIED is a list of components that were examined during the execution of the module 202 in regard to the suitability thereof for resolving a specific processed action and were found to be inappropriate. TRIED is built and examined in relation to each action to be resolved in order to prevent endless loops. The output of module 202 consists of a partial or a complete DONE list and a data element holding a Boolean value specified as RESULT which

indicates the outcome of module 202 performance. The two predefined values RESULT indicator could be set to are "success" or "failure".

At step 206, the ACTIONS list is examined in order to check whether the list still stores any unresolved actions. If ACTIONS is empty at step 208, the
5 DONE list is the final output. At step 209 RESULT indicator is set to the value of "success" thereby enabling at step 210 the return of RESULT indicator to the calling routine indicating success.

If at step 206 ACTIONS list found to hold unresolved actions then at step 212 the first action in ACTIONS, specified as action A, is copied to a
10 location of routine held in allocated physical location on the storage device 12. At step 214 action A is examined in regard to the type thereof. If the type of A found to be a trial action then program control is passed at step 216 to a sub-routine of module 202 for the handling of trial actions. Said sub-routine will be described hereunder in association with the following drawings. If the type of A was
15 checked and found not to be a trial action then at step 218 the sub-routine for creating trials is called. The create trials sub-routine will be described hereunder in association with the following drawings.

The execution of the trial-creating sub-routine effects the creation of a list of trial actions TRIAL-LIST. TRIAL-LIST is built by analyzing action A,
20 extracting the operative steps from action A, and inserting said steps as trial actions into TRIAL-LIST. In effect, trial actions in TRIAL-LIST replace action A. Therefore at step 219, action A is removed from the ACTIONS list. At step 220, the first trial action from the list of trial actions TRIAL-LIST, specified as T, is copied to a working segment of the routine. At step 222 it is examined whether
25 T holds a valid trial-action. If T is a valid trial action then a new action list ACTIONS is created by the insertion of T into ACTIONS at step 226. Consequently, the module 202 is called recursively with the group of the predefined parameters, which are set to different values for each recursive call. The called module 202 returns with a new DONE list having different actions stored,
30 and a new RESULT indicator storing the values of either "success" or "failure".

Subsequently, program control passes back to step 220 in order to handle the next trial action T from TRIAL-LIST. If no more trial actions are found in TRIAL-LIST at step 222, then RESULT indicator is examined at step 223 whether the prior recursive call produced a value of "success" or "failure".

5 According to the value of RESULT an additional indicator specified as TOTAL-SUCCESS, is set to the value of "success" or "failure" at step 225 and returned to the calling routine.

Referring now to Fig. 4 that illustrates the execution flow of the
10 routine that handles non-trial actions. At step 260 action, A is checked for the relevant xor rules by calling the appropriate subroutine and as a result, an action B is created. The detailed description of the called subroutine at step 260 will be set forth hereunder with association with the following drawings. At step 262 a check is performed on action B to determine whether the action is having any valid
15 executable steps such as installing or uninstalling components. If there are no such steps, at step 266 action A is removed from the list of ACTIONS and at step 268 the Resolve Module 202 called recursively with the new ACTIONS. At step 286 the module returns with the appropriate value stored in RESULT indicator.

If there are valid executable steps in action B then at step 264 B is
20 examined whether the executable steps involve installation of components. If no such steps are found then at step 284, program control passes to step 311 of Fig. 5. The detailed description of the flow logic therein will be set forth hereunder with the associated drawing.

When executable steps involving installation of components are found
25 in action B at step 264 then at step 270 the component to install is copied to a working segment of the routine specified as X. Subsequently, at step 272 a check is made whether the list DONE contains an action to uninstall component X. If such an action found then the module 202 returns the value "failure" indicative of the failure of resolving action A. If the result of the check is negative then next at
30 step 274, the TRIED list is checked regarding the existence of X therein. If X not

exists in the TRIED list then at step 282 X is added to TRIED, and program control passes to the add-remove rules checking routine. If TRIED does not contain component X then at step 276 ACTIONS is scanned to search for a trial action that installs X. If such an action was not found then program control passes
5 to step 311 of Fig. 5 for further processing. In contrast, when a trial action that installs X was found, then at step 280, A is removed from ACTIONS, at step 281 B is added to DONE, and at step 268 the Resolve Module is called with the new set of input parameters. The Resolve Module 202 returns an indicator as a result, specifying "success" or "failure".

10

Referring now to Fig. 5 showing the continuation of the flow of control from Fig. 4. At step 311 the check-step module is called with B passed as parameter and RESULT as the returned value. The operation of the module will be described hereunder in association with the accompanying drawing Fig. 10.
15 The value of RESULT is examined at step 325. If the RESULT indicator returned by the module is 'failure' then at step 315 the RESULT is returned to the calling routine. If the value stored in RESULT is "success" then at step 313 the check-add-remove rules module is called. The operation of the module will be described hereunder in association with the accompanying drawing Fig. 9. At step 312 the
20 PRE-ACTION list is examined related to the existence therein of any actions. If the list is empty then action B is appended to the DONE list at step 322, action A is removed from ACTIONS list at step 324, POST-ACTION list is appended to ACTIONS list at step 326 and the resolve module 202 is called recursively at step 318. At step 320 after program control returns from the called routine, RESULT
25 indicator is returned to the calling routine.

If at step 312 PRE-ACTION list is found empty then steps 314 through 320 are executed. At step 314 A is removed from ACTIONS and at step 316 the ACTIONS list is re-built by appending PRE-ACTION list, A, POST-ACTION list, and ACTIONS into ACTIONS. At 318 the resolve module 202 is called

recursively. When program control returns to 320 the RESULT indicator returned from the called routine is returned to the calling routine

Referring now to Fig. 6 that illustrates the execution flow of the xor-
5 rules-checking routine generally referred to as module 282. In frame 284 the input and the output data elements associated with the routine are listed. The input to the routine consists of A, which is a trial action passed from step 260 of Fig. 4, and S the relative system. The relative system is defined as the pairing of the system and the actions already resolved and to be performed on the system at any
10 specific point in time during the execution of the module. The relative system S consists of the list of components installed in the system and a list of actions to perform thereon. In order to examine whether a component is in the system the relative system should be tested by a) checking if the component is in the system and b) examining the list of actions for the component in order to determine a
15 potential installation or a potential removal of the component. The output of the module 282 is B, an action created as a result of the xor-checking of action A.

At step 286 the relative system S is tested to determine whether S contains the action for the installation of A specified as A.TO_INSTALL. If the result is positive then A.TO_INSTALL is copied to the B.X variable at step 288.
20 At step 290 the xor set of the B.X is extracted from the knowledge base 28 of Fig. 1 by utilizing the xor rules table 32 of Fig. 1 respective to the action B.X. As a result at step 292 a set of actions U is created. U will consist of an A.TO_UNINSTALL action and the xor set thereof. At step 294 a loop control variable INDEX is set up and initialized to control the execution of the loop
25 across step 296 through step 304. The loop is performed in order to examine whether S, the relative system, contains any of the actions U consists of. At step 296 it is determined whether all the actions in U were examined. If the result is positive the loop terminates, the value of U is copied to B at step 297 and the value of B is returned to the calling routine at step 298. If not all actions in U
30 were examined then the loop continues to execute. At step 300 S is scanned to test

if the specific action in U, specified as U[INDEX] exists in S. According to the result of the test either U[INDEX] is removed from U at step 302 and the loop continues to execute until all of U is examined or the loop control variable INDEX is incremented at step 304 and the loop continues to execute until all actions in U were examined. If at step 286 the relative system S is found not to contain A.TO_INSTALL then at step 306 B.X is initialized to null, at step 308 A.TO_INSTALL is copied to U, and program control passes to step 294 in order to begin the loop across steps 296 and 304.

Referring now to Fig. 7 that illustrates the flow of logic for the create-trials-for-install-action module, called at step 218 of Fig. 2. The module contains three functions: a) below(N) that accesses the component data table 28 of Fig. 1 with node N and produces an ordered set of components below the node N, (inclusively, i.e., N is below N), b) element-of(N) that accesses the component data table 28 of Fig. 1 with node N and returns the element E above N, and c) contains(N) that accesses the relative system S with node N and returns indication whether S contains N, i.e., whether S contains any of below(N).

As shown at 133, the input for the create-trials-for-an-install-action module 131 consists of A and S. A is an install action comprising a list of nodes. The number of resulting trial actions will be equal to the total number of components below the nodes. S is the relative system comprising the system and the list of components to install at any given point of time during the execution of the resolve module 202. The output of the module 131 is the 'TRIAL-LIST' consisting of all the trial-actions resulting from the input A. At step 130 two list variables L1 and L2 are initialized to null. At step 132 a loop across the nodes of A is activated. At step 134 A is examined whether there are any more nodes therein. If there are nodes to process therein then at step 142 the next node in A is copied to a physical location allocated on the storage device 12 that is specified as working segment N. Consequently, at step 144 the function element-of(N) is called in order to obtain from the component data table 28 of Fig. 1 the element

above N. At step 144 a test is made to determine if the element-of(N) is N itself. If the result is positive at step 150 the below(N) function is called and resulting set of components is appended to the list L1. The loop continues at step 134. If the result of step 144 is negative then at step 146 the contains(N) function is activated
5 in order to determine whether the relative system S contains N. When the result positive step 150 is executed by appending the set of components obtained to the list L1. Negative result effects the execution of step 148 in which the set of components obtained from function below(N) is appended the list L2. Both conditions result in the continuation of the loop at step 134 until all the nodes in A
10 are processed. When the required operations on A are completed, the loop terminates and the steps 136, 138, and 140 are performed to create TRIAL-LIST by processing lists L1 and L2 respectively. For each component X on the lists L1 and L2 a trial action "install X" is created, and appended to TRIAL-LIST. At step 140 the entire TRIAL-LIST is returned to the calling module.

15

Referring now to Fig. 8 that illustrates the flow of logic for the creation-of-trials-for-update-action module called at step 218 of Fig. 2. As shown at 163 the input for the module 161 is X, a component, and the output is the TRIAL-LIST. At step 165 the list specified as CANDIDATES is created by
20 calling the functions element-of(X) and below(result of element-of(X)). At step 160 a loop is activated on the CANDIDATES. The exit from the loop is controlled at step 162 by examining whether there are any unprocessed elements left in the list CANDIDATES. At step 166 an element from the list is copied to a working segment specified as C and examined at 168 in order to determine
25 whether C is equal to the component X. If C is X then at 170 the trial actions "install C" and "uninstall X" are created and appended to TRIAL-LIST. The loop continues at step 162. When all the elements in CANDIDATES were processed, an exit from the loop is effected at step 162 and TRIAL-LIST is returned to the calling routine at step 164.

30

Referring now to Fig. 3 that illustrates the steps of the 'relative-system-contains-X' module. The module is called from step 146 of Fig. 8 in order to examine whether the relative system contains the node X. The input and the output data elements of the module are shown in frame 231. The input comprises
5 of S the relative system, and X the node to examine. The output is specified as H, which is an indicator the value of which is to be set to "false" or "true" consequent to the module operation. At step 232 the size of DONE is calculated and copied to a physical location allocated on storage device 12, intended to be used as working segment and specified as L. At step 234 a loop control variable
10 INDEX is initialized and the loop is activated at step 238. When the examination of the entire contents of DONE is completed, the current value of H is returned to the calling module. As long as the loop is active the following steps are executed for each component in the list DONE: At step 242 the value of H is tested. If the value is "true" then at step 244 a check is made to determine whether X is equal to
15 the uninstall action in DONE which is specified as DONE[INDEX].UNINSTALL. When the result of the compare operation is positive H is set to "false" at step 248, INDEX is incremented at step 252 and program control returns to step 238 in order to continue the execution of the loop. When the result of the compare operation is negative, at step 252 INDEX is
20 incremented and the loop continues to be performed. If at step 242 the value of H is not "true" then at step 246 a check is made to determine whether X is equal to install action in DONE specified as DONE[INDEX].INSTALL. A positive result will effect the setting of the value of H to "true" at step 250. Subsequently at step 252 INDEX is incremented and control returns to 238 to continue the execution of
25 the loop. A negative value will effect the execution of step 252 only without setting the value of H and return to 238 to continue the execution of the loop.

Referring now to Fig. 9 which illustrates the flow of logic of the check-add-remove rules module 172. The input and the output data elements of module
30 172 that are shown in frame 173. The input consists of S the relative system, X the

node to install, and C the list of nodes to uninstall. The output of module 172 consists of the lists PRE-ACTION and POST-ACTION.

Module 172 updates S at step 174 by adding node X and removing C from DONE. At step 174 two sets of rules are created by accessing the knowledge
5 base 26 of Fig. 1 and activating the appropriate routines. The set of rules relating to the installation of X is specified as R1 and the set of rules relating to the uninstalling of C is specified as R2. At step 178 the first loop is activated in order to process R1. The loop control is operative at step 180. When all the processing of the rules in R1 completed, the first loop is terminates and control passes to 190
10 to activate the second loop to handle R2. The first loop consists of steps 182 through 188. At step 182 the next rule in R1 is copied to physical location allocated on the storage device 12, intended to be used as a working segment, and specified as R. At step 184 the "relative-system-contains-X" module is called to examine whether S contains components associated with the rules related to X. If
15 the result is affirmative, control passes to step 186 to examine whether S contains components associated with rules related to C. If the result of the test is negative an install action "install R.NEED" is created at step 188 and appended to PRE-ACTION list. In contrast when the result of the test is negative program control passes to step 180 to continue the execution of the first loop. If the result of the
20 operation at step 184 is negative then program control passes back to step 180 to continue the execution of the loop.

When the first loop is completed the second loop is activated at step 190 to process list R2. The loop control is operative at step 192. When all the processing of the rules in R2 is completed, the second loop terminates and control
25 returns to the calling routine. The second loop consists of steps 192 through 201. At step 194 the next rule in R2 is copied to a physical location allocated on the storage device 12 intended to be used as a working segment, and specified as R. At step 196 the "relative-system-contains-X" module is called to examine whether S contains components associated with the rules related to X. If the result is
30 affirmative, program control passes to step 198 to examine whether S contains

components associated with rules related to C. If the result of the test is negative then at step 200 an update action "update node of C" is created and added at step 201 to the POST-ACTION list. In contrast when the result of the test is positive, program control passes to step 192 in order to continue the execution of the second loop. If the result of the operation at step 196 is negative then program control passes back to step 192 to continue the execution of the loop.

Referring now to Fig. 10 that illustrates the flow of logic of the check-step module 80. The module 80 associated with the processing of the locks, the user's preferences and the mode of operation of the virtual upgrade process. The input and output data elements of the module 80 are shown in frame 81. The input consists of the following parameters: MODE specifying the method of operation such as "interactive mode", B the action to perform, LOCKS a set of components that were locked by the user in order to present installation or removal thereof, and NOBACK a set of elements indicating the set of versions for a specific component. The output is RESULT that would be set to the value of "success" or "failure" according to the result of the operation. The module 80 is called from step 311 of Fig. 5. B is the action to perform and consists of B.X that are components to install and B.CI that are components to uninstall.

At step 82 the first element of B.X is copied to physical location allocated on the storage device 12, intended to be used as a working segment, and specified as E and at step 84 a loop is initiated over B.CI. The loop control is operative at step 86. As long as the loop is active the steps 86 through 112 are executed. At step 90 LOCKS is examined to check whether LOCKS contains B.CI. If the result of step 90 is positive at step 92 MODE is examined to check whether the value thereof is "interactive mode". If the result is negative, RESULT is returned to the calling routine with the value of "failure". If the result of step 92 is negative at step 96 the user's permission is requested to remove the lock from the component. Subsequent to the user's refusal at step 100 RESULT is set to the value of "failure" and returned to the calling routine. In contrast, if the user allows the removal of the lock from the component at step 100 a series of tests are made

at steps 98, 104, and 106 respectively. At step 98 NOBACK is examined in regard to the presence of B.CI therein. If B.CI not in NOBACK control returns to step 86 to continue the loop. If NOBACK contains B.CI, at step 106 the location of B.CI relative to E and B.X is tested. If B.CI is below C and before B.X in the component data 28 of Fig. 1 then MODE is checked regarding the method of operation. If the method is not interactive RESULT is set to "failure" and program control returns to the calling routine at step 108. In interactive mode the user is asked for permission to go back in component versions. If permission is not given the value of RESULT is set to "failure" and returned to the calling program at step 114. If user permission was granted, program control returns to step 86 to continue with the execution of the loop. Consequent to the successful processing of the entire set of components to uninstall the loop terminates at step 86, the value of RESULT is set to "success" and at step 88 RESULT IS returned to the calling routine.

15

The above described resolve module and the associated functions thereof are proposed as central constituent of an integrated tool intended to resolve inter-component dependencies among hardware and software entities in an operative computer system and thereby accomplish a seamless system upgrade. For each hardware and software entity to be added to the system or removed from the system the resolve module analyzes the entity compatibility requirements in regard to other related software, hardware or operating system entities. The analyzing consists of locating and processing the entire set of related entities in accordance with a set of predefined rules. As the process is recursive by nature, i.e., each new entity processed and inserted into the system following of the requirements of the original entity, necessitates further analysis in regard to still another set of entities related to the newly inserted entity, further processing, and further insertions into the system, where each new action involves a new set of actions. After all the entities are processed, practically all dependencies are

20
25

resolved and the system contains only entities that are capable functioning together as an operative system.

The present invention proposes to include a number of useful additions and substantial modifications to the method disclosed by the above-referred patent applications. The additional software elements, rules, sub-methods, and data
5 structures constituting the proposed method will be described next:

ADDING COST TO CHANGES

Conflict Resolver module 40 outputs a set of lists concerning recommended upgrades to the original system. The lists are compared in order to
10 determine installation preference i.e., to decide which list of upgrades from the lists-set is preferable to the other lists. In order to receive a precise solution a "cost" or a number is assigned to each upgrade (or "change"). The cost is then summed up to provide a total cost. The cost of each change may depend on a number of parameters such as: a) upgrade class (software, kernel parameter,
15 module), b) upgrade route (upgrade to newer version or upgrade to older version), c) upgrade-associated-process (compilation or no compilation), d) upgrade type (installation or un-installation), and the like.

Cost is actually a "price" label predefined by the user or the system administrator. In the preferred embodiment of the present invention the price label
20 is set as the period of time necessary for the performance of the specific upgrade regarding the selected component in the original system. In order to provide a thorough understanding of the cost-related or "price"-related selection of a list out of a set of lists a number of examples will be given next.

Example 1: A costly action could be a "compilation". A compilation is
25 a process, which will be invoked following the installation of a new kernel. Subsequent to the action of installing the new kernel the kernel should be compiled. As the compilation of the kernel involves a relatively "long" period of time the original action will carry the price tag of a 1000.

Example 2: The installation of "patch" also requires a compilation and
30 various adjustments (such as the updating of system tables). Therefore the action

will be given a price tag of 1200. The above cost is combined from the price of the compilation, which is 1000 and the price of the adjustments that could be 200.

Example 3: The installation of a kernel parameter also involves compilation and various adjustments. The price could be fixed as 1100. The price
5 is summed from "compilation" with a value of 1000 and from "parameter adjustment" (such as updating of other system tables) with a value of 100.

Example 4: As the installation of a software program demands a relatively shorter time period the cost of the action could be set as 500.

Although each action has a specific price tag the summation
10 of the action does not necessarily involve an arithmetic summation of all the price labels. A compilation is made only once for all the actions. Therefore when a new kernel is installed and on the top of the kernel new patches are installed the total cost for indicated for compilations will be only a 1000.

The Cost table 37 of Fig. 1 is operative in storing the list of
15 cost-related parameters with the associated "cost" value thereof. Note should be taken that a "cost" value could be associated with several parameters.

HANDLER RULES

The add-remove rules stored in the add remove rules table 34 of Fig. 1 and xor-rules stored in the xor-rules table 32 of Fig. 1 are utilized in order to
20 determine the correct dependency relationships among a set of potentially installed or uninstalled components. Thus the add-remove rules and the xor-rules are utilized to assure the correct operation of an upgraded component in association of with other components in the system or the correct operation of components in association with the lack of a previously installed component.
25 However, there are some rules that are relevant only during the installation process of a specific component. For example the kernel needs a compiler to be installed but after the installation the compiler is not needed. In order to deal with similar situations "handler rules" are utilized. The handler rules are stored in the handler rules table 35 of Fig. 1. Handler rules are utilized at the resolved stage of
30 the process and do not participate in the checking process. Therefore handler rules

are attached to the component's installer (handler) script. The handler rules are activated only at that point of time when a decision is taken to execute the appropriate upgrade involving the specific component. There are two types of handler rules: a) common rules and b) specific rules. Common rules concern a set of components while specific rules apply to particular components.

One possible example for a common rule could be:

"source-rpm_handler NEEDS C++ compiler"

The above rule determines that for the installation of each RPM component in source code format the presence of a C++ compiler is required in the system. When the specific component is about to be installed the component's handler (or installer routine) examines the handler rules table 35 of Fig. 1 regarding handler rules associated with the specific component. When an associated handler rule is found the rule is analyzed and the original system is checked for the presence of the required compiler. If not found the compiler is downloaded from the central server and appropriately installed. One possible example for a specific rule could be:

"some_software NEEDS perl"

The above rule determines that for the installation of a component containing some software the presence of the Perl interpreter is required in the system. When the specific component is about to be installed the component's handler (or installer routine) examines the handler rules table 35 of Fig. 1 regarding handler rules associated with the specific component. When an associated handler rule is found the rule is analyzed and the original system is checked for the presence of the required interpreter. If not found the interpreter is downloaded from the central server and appropriately installed.

WORKING RULES

Working rules define specific common-sense working habits such as the order of installation for components of different types. One such a working rule may define that a kernel patch shall be installed after the installation of full kernel or that a software module shall be installed after the installation of a full

kernel. In the preferred embodiment of the present invention working rules are activated at that point in time when the real upgrade of the components on the selected action list is initiated.

NON-BINARY CHANGES

5 The resolve decision in the related patent applications is based on binary data only representing various flags returned from the resolver modules. All the various operative routines return indications as values represented by one or more binary digits signifying parameters such as "success"/"failure", "exists/not exists" and the like. In the proposed improvements to the present
10 invention data in other format could be returned e.g., the size of the RAM, the port number used by IRC, and the like. Non-binary data returned from the specific routines could be handled by a generalization of the "rule" and "change" concept.

RESOLVE FAILURE REASONS

15 The method disclosed in the related patent applications involved the return of a "success/failure" flag only in order to indicate success or failure regarding the process of resolving inter-componental dependencies. After failure the proposed improved method will return, in addition to the indicator flag, appropriate error messages concerning the reasons for the failure. The error messages will be then displayed to the user/system administrator. Thus the users
20 will have the option of modifying their requests concerning the number or type of components to be installed or uninstalled.

TREE SEARCH METHODS

25 In order to optimize the search process a number of improved tree-search techniques are utilized in place of the recursive search used in the related patent applications. Some of the techniques are used in combination with others to achieve the shortest possible vtree search times. The algorithms utilized are known in the art. In the preferred embodiment of the present invention the search algorithms utilized are a simplified Dijkstra's shortest path algorithm, a variation of the greedy search algorithm and a combination thereof. The combined
30 algorithm is a modified fixed-depth tree search through which the entire vtree is

scanned to a certain depth, the optimal branch is selected, and the scan is continued from it onwards slightly deeper. Selecting appropriate algorithms to be used in specific routines accomplishes substantially faster processing.

KNOWLEDGE IN THE BRANCH

5 If an acceptable solution exists for resolving dependencies the use of the proposed search algorithms substantially decrease the period of the processing. However, if no solution found then the entire tree must be scanned. By storing "Knowledge" in the branches of the tree the effective size of the tree can be reduced. The reduction can be accomplished by keeping knowledge of rules that are satisfied in a branch of the tree. If in an exemplary branch of the tree the action "install A" and in the rules table a rule exists regarding the installation of component A such as "A NEEDS B" the simplified rule such as "B NEEDED" is stored in the particular branch. Therefore any further request to uninstall B will cause the branch to fall immediately.

ESTIMATED COST OF BRANCHES

15 The size of the search tree is substantial therefore the searches are long. The cost of a branch in the tree depends on the actual changes done in the branch up to a certain point in time and not on the actions yet-to-be-done. The actions' cost can not be precisely calculated but it can be estimated. Consequently the branch that has the least estimated cost will be selected earlier. The cost value of the actual changes done is stored in the Cost table 37 of Fig. 1.

FIXING PRIORITY

After executing the method disclosed by the related patent applications the result received was a list arranged in the correct order of dependencies, i.e., if 25 A needs B then B should be installed before A. The method proposed by the present invention outputs the list in a different order, and therefore the output should be sorted. The sort is performed according to the dependencies, and according to the component types. Thus after sorting the input correctly an exemplary kernel base component will appear before the kernel patches thereof.

Persons skilled in the art will appreciate that the present invention is not limited to what has been particularly shown and described hereinabove. Rather the scope of the present invention is defined only by the claims, which follow.

CLAIMS

I/We Claim:

1. Within a computer system having a central processing unit, a storage device, a communication device, a method of assisting users of computer client systems in the conduct of a requested computer system upgrade process by automatically establishing an operatively correct combination of a set of components in the computing environment, comprising:
5 obtaining control tables from a storage device, having information necessary for the execution of said method;
10 setting up result tables on said storage device, to hold the information resulting from said establishing process;
examining said control tables for dependency conflicts arising among said components in control tables;
identifying said dependency conflicts arising among said components in
15 said control tables;
resolving said dependency conflicts among said components in said control tables, thereby creating result-related actions.
2. The method of claim 1 of obtaining control tables further comprising
20 the steps of:
obtaining a system information table comprising a set of operative components installed in a computer system;
obtaining a first system upgrade table comprising a required set of operative components to be installed in said computer system;
25 obtaining a component objects table comprising a set of installation-related components;
obtaining a component database that contains a structured index of said component objects table;
obtaining a rules table comprising installation-related dependency rules
30 referring to said components.

3. The method of claim 1 of setting up tables further comprising the steps of:

building a resolved system upgrade actions table for the purpose of storing
5 dependency conflict-free system upgrade actions;

building a second system upgrade table for the purpose of storing
unexamined system upgrade actions;

building a third system upgrade table;

10 copying said first system upgrade table into said second system upgrade
table.

4. The method of claim 1 of resolving further comprising the steps of:

obtaining from said second system upgrade table a system upgrade action;

15 disassembling said system upgrade action obtained from said second
system upgrade table into at least one action containing at least one

component and at least one upgrade operation;

inserting said set of at least one action into said third system upgrade table;

obtaining said at least one action from said third system upgrade table;

checking xor rules for said at least one action;

20 handling user preferences concerning the execution of each member of said

at least one action arising from dependency resolving;

checking add-remove rules for said at least one action;

inserting resulting said at least one action to said resolved system upgrade
actions table;

25 removing said at least one action from said second system upgrade table

thereby indicating that said system upgrade action is dependency conflict-
free.

5. The method of claim 4 of checking xor rules further comprising the steps of:

examining the system information table in combination with at least one

5 resolved system upgrade action corresponding to at least one action;

obtaining at least one xor rule concerning said at least one action from said xor rules table;

combining said xor rules with said at least one action;

10 checking said combined set of at least one rule with at least one action in regard to existence in said system information table;

checking said combined set of at least one rule with at least one action in regard to existence in said resolved system upgrade actions table;

removing actions from said combined set of at least one rule with at least one action which do not appear in said system information table and said

15 resolved system upgrade actions table.

6. The method of claim 4 of checking add-remove rules further comprising the steps of:

20 obtaining said at least one action-related add-remove rule from said add-remove rules table;

resolving the added install actions required by said at least one action;

creating a list of at least one install action corresponding to said add-remove rules;

resolving the added replace actions required by said at least one action;

25 creating a list of at least one update action corresponding to said add-remove rules.

7. The method of claim 4 of handling user preferences further comprising the steps of:

obtaining at least one user preference record relating to at least one lock placed on at least one action;
obtaining at least one user preference record containing a user authorization mechanism concerning the utilization of prior versions ;
5 determining the mode of operation;
examining said at least one action in association with at least one user preference record and said mode of operation;
determining whether said at least one user preference record is an impediment to said resolving process;
10 enabling manual removal of at least one user preference record .

8. The method as claimed in claim 7 where said user preference record is enabled to be removed is an impediment to the conflict resolving process.

15 9. The method of claim 1 further comprises inserting resolve-related actions into said result tables.

10. The method of claim 1 further comprising notifying said user of said computer system of resolve-related actions thereby accomplishing a real
20 system upgrade by outputting content of said result tables.

11. The method as claimed in claim 1 where said system upgrade action comprises at least one component and at least one upgrade instruction.

25 12. The method as claimed in claim 1 where said set of components comprising hardware components, software components and system components.

13. The method as claimed in claim 1 where said result-related actions are
30 appended to said system upgrade actions and processed like the same.

14. The method as claimed in claim 2 where said component database is organized in a tree-like hierarchical data structure.

5 15. The method as claimed in claim 2 where said component database contains component indices connected according to the explicit dependency-relationships among said components.

10 16. The method as claimed in claim 2 where said rules in said rules table are encoded into a specific formal logic-like format.

17. The method as claimed in claim 3 where said component database is linked to said rules table by said component indices.

15 18. The method as claimed in claim 3 where said second system upgrade table is a temporary table.

19. The method as claimed in claim 3 where said second system upgrade table contains unexamined system upgrade actions.

20 20. The method as claimed in claim 3 where said third system upgrade table is a temporary table.

25 21. The method as claimed in claim 1 wherein the installation preferences are determined according to pre-determined values assigned to the recommended upgrade actions.

30 22. The method as claimed in claim 1 wherein installation-specific handler rules are utilized to determine the necessity for the presence of specific software modules for the duration of the installation.

23. The method as claimed in claim 1 wherein specific working rules are utilized to determine the order of installation for components having different type.

5

24. The method as claimed in claim 1 wherein the resolver modules provide non-binary parameters to the associated modules in order to improve on the inter-module communication.

10

25. The method of claim 1 wherein in addition to the indication concerning the success/failure of a specific operation the resolver modules provide additional information in order to improve on the inter-module communication.

15

26. The method of claim 1 wherein the combined tree search methods are utilized in order to improve on the running time of the method.

20

27. The method of claim 1 wherein specific information concerning the optimal installation process is stored within the structure of the vtree in order to improve on the execution time of the method.

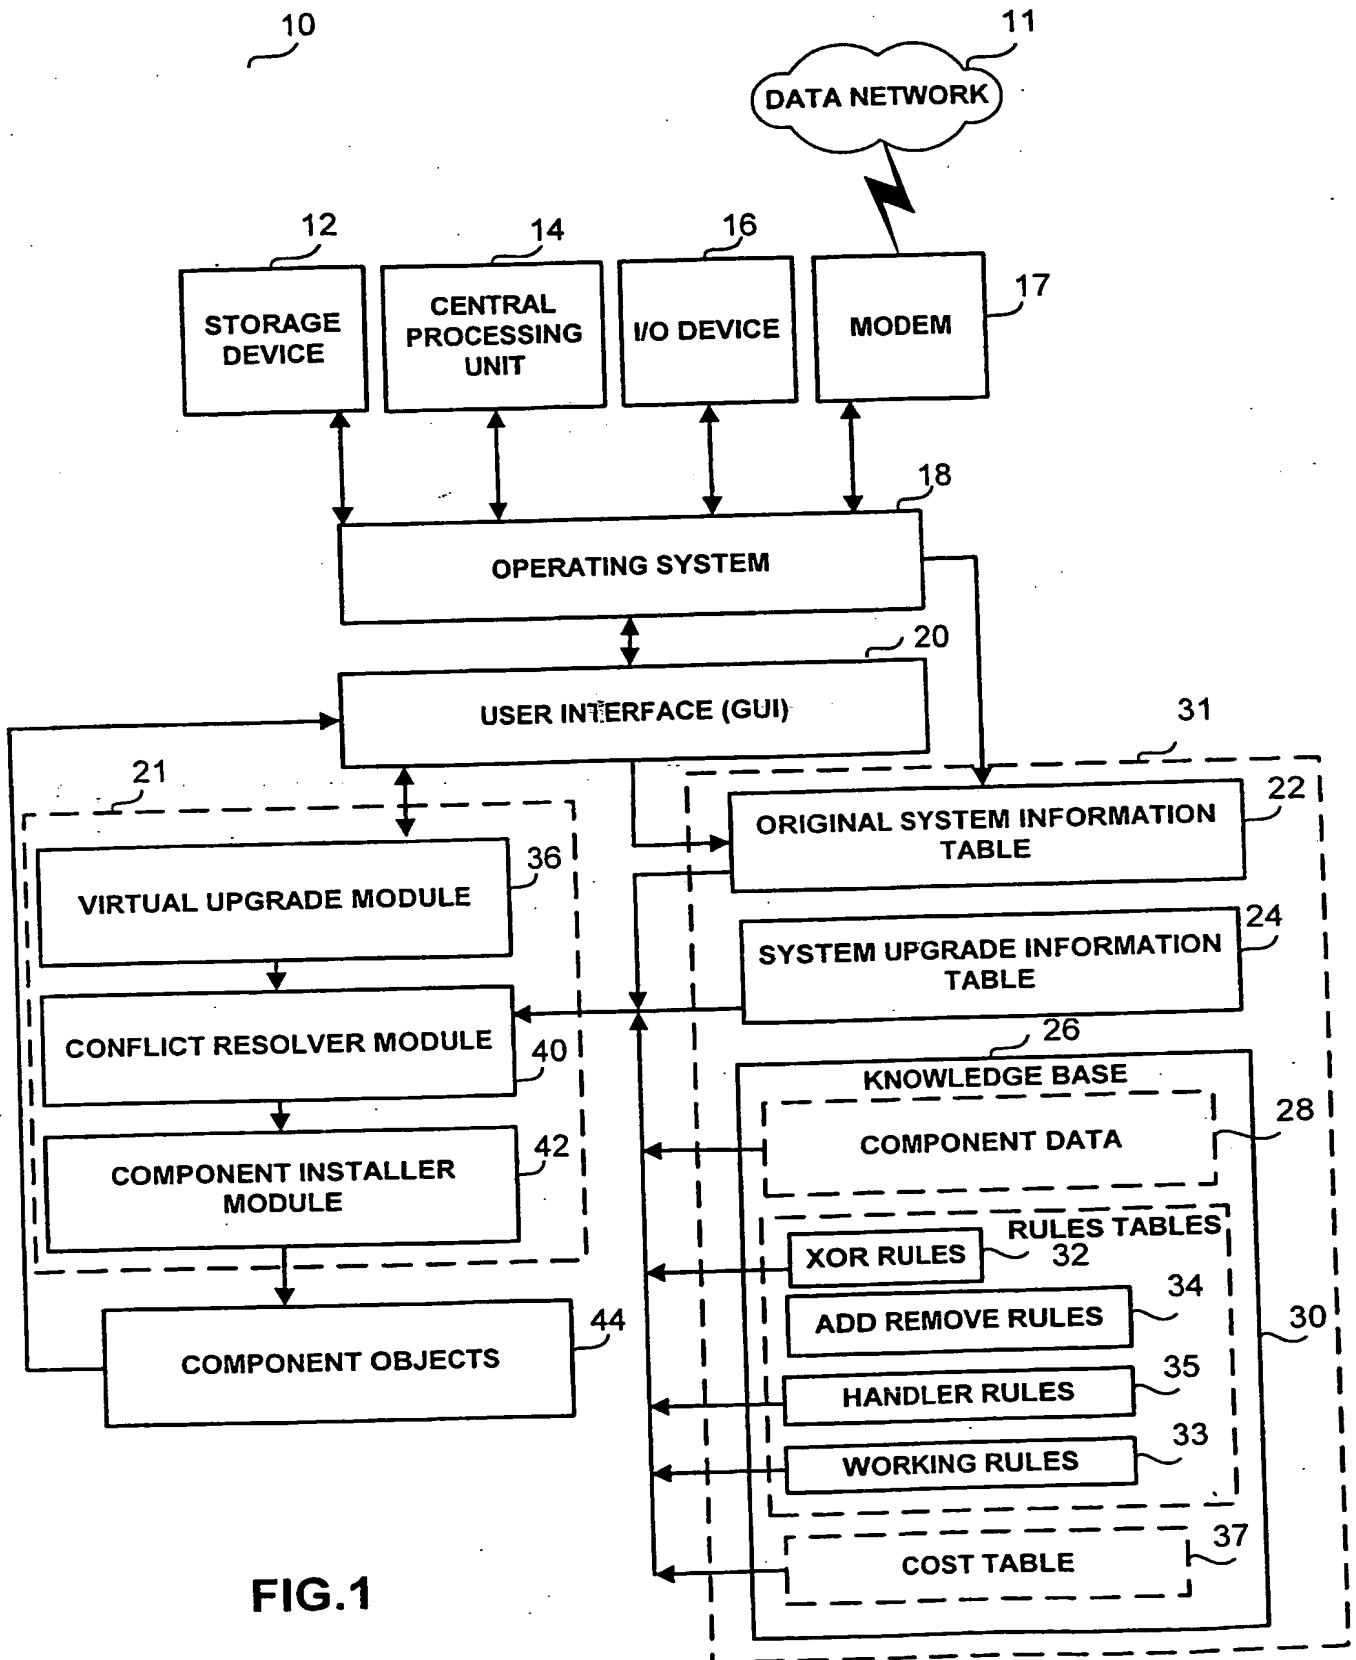
25

28. The method of claim 1 wherein the installation preferences determined according to pre-determined values assigned to the recommended upgrade actions are tree-structure specific.

29. The method of claim 1 wherein the result list is ordered in accordance with the dependencies and the component type.

30

1/10



2/10

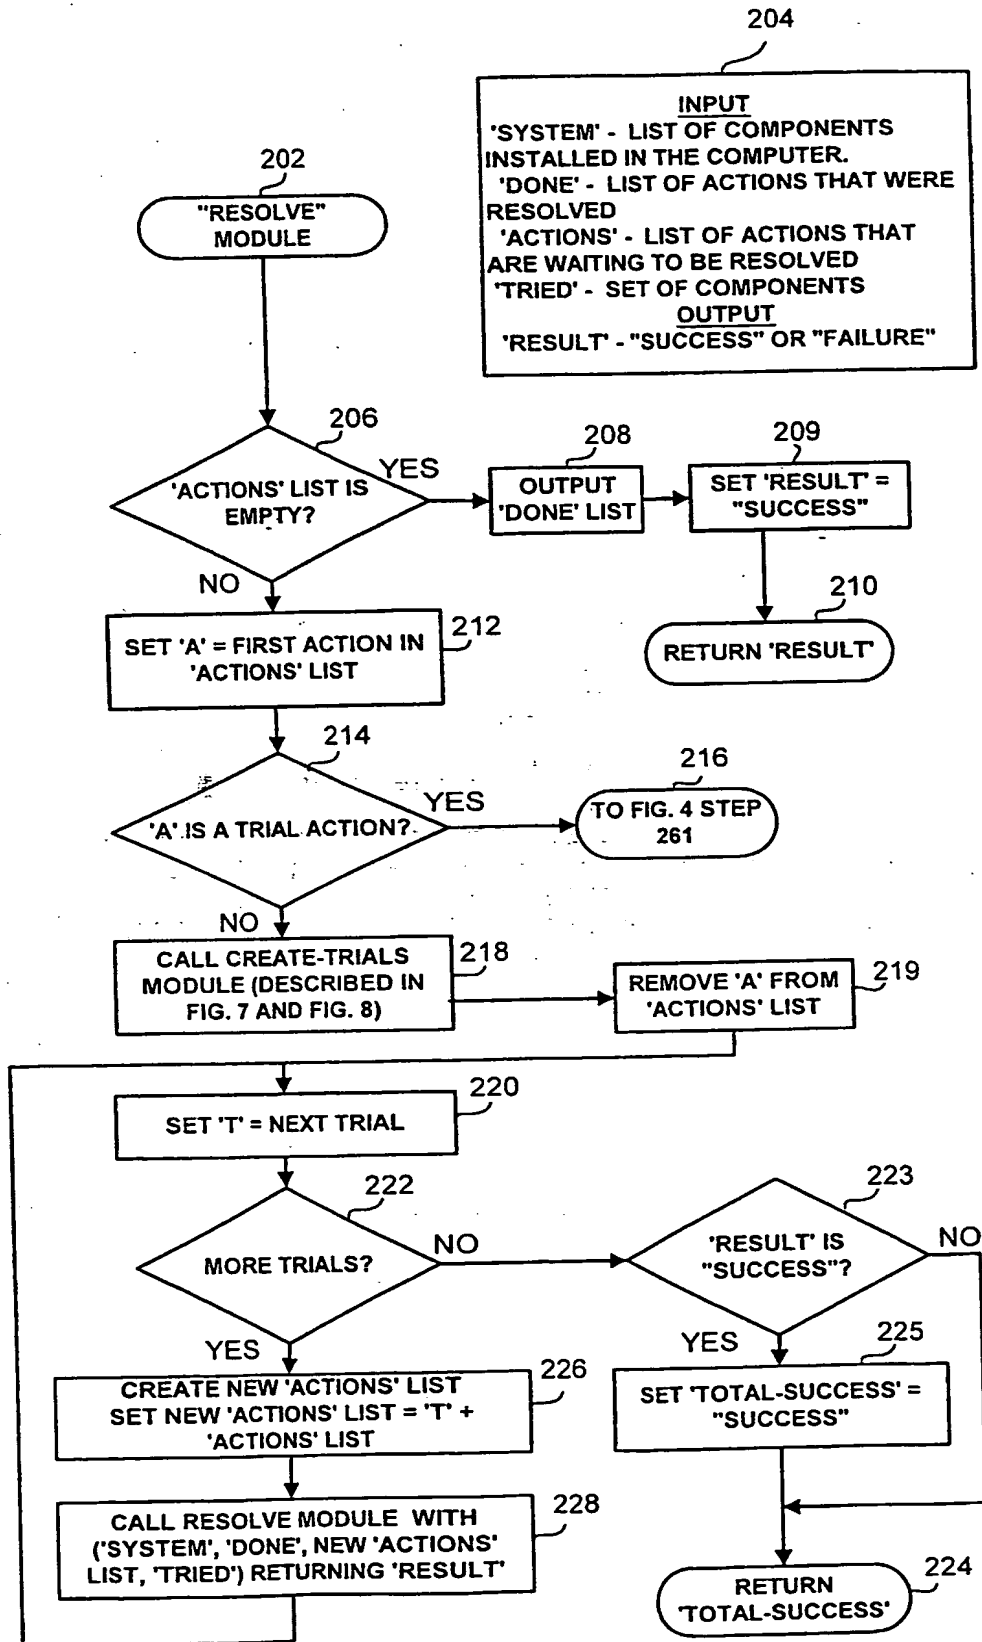


FIG. 2

3/10

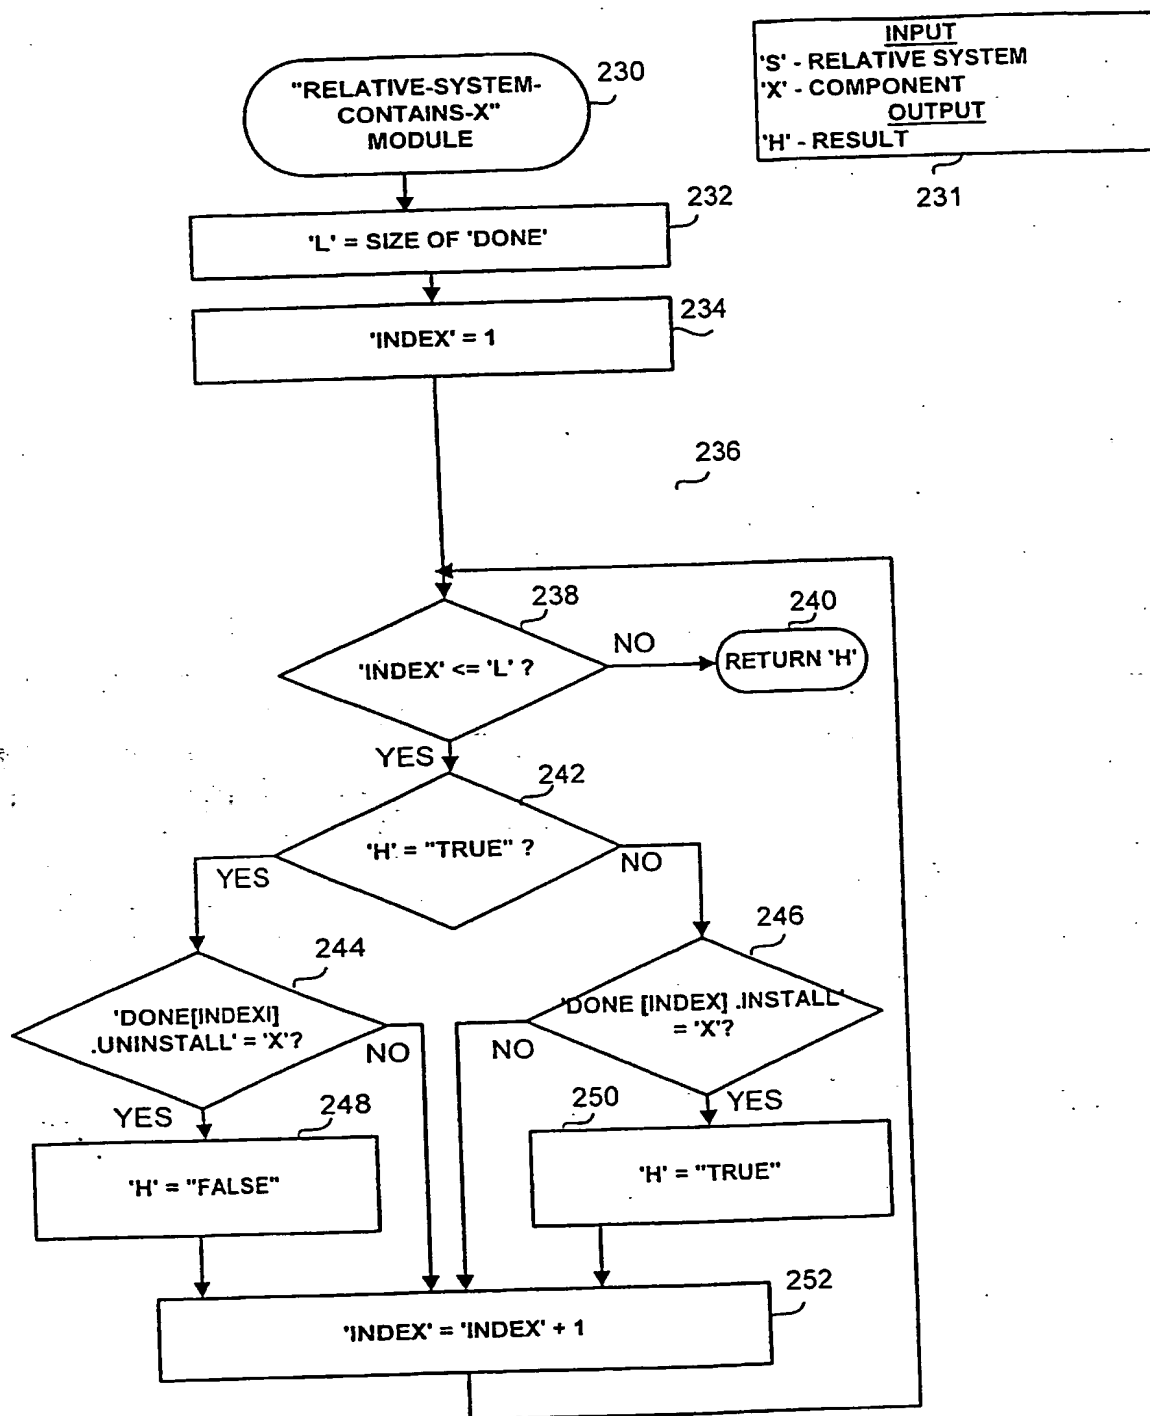


FIG. 3

4/10

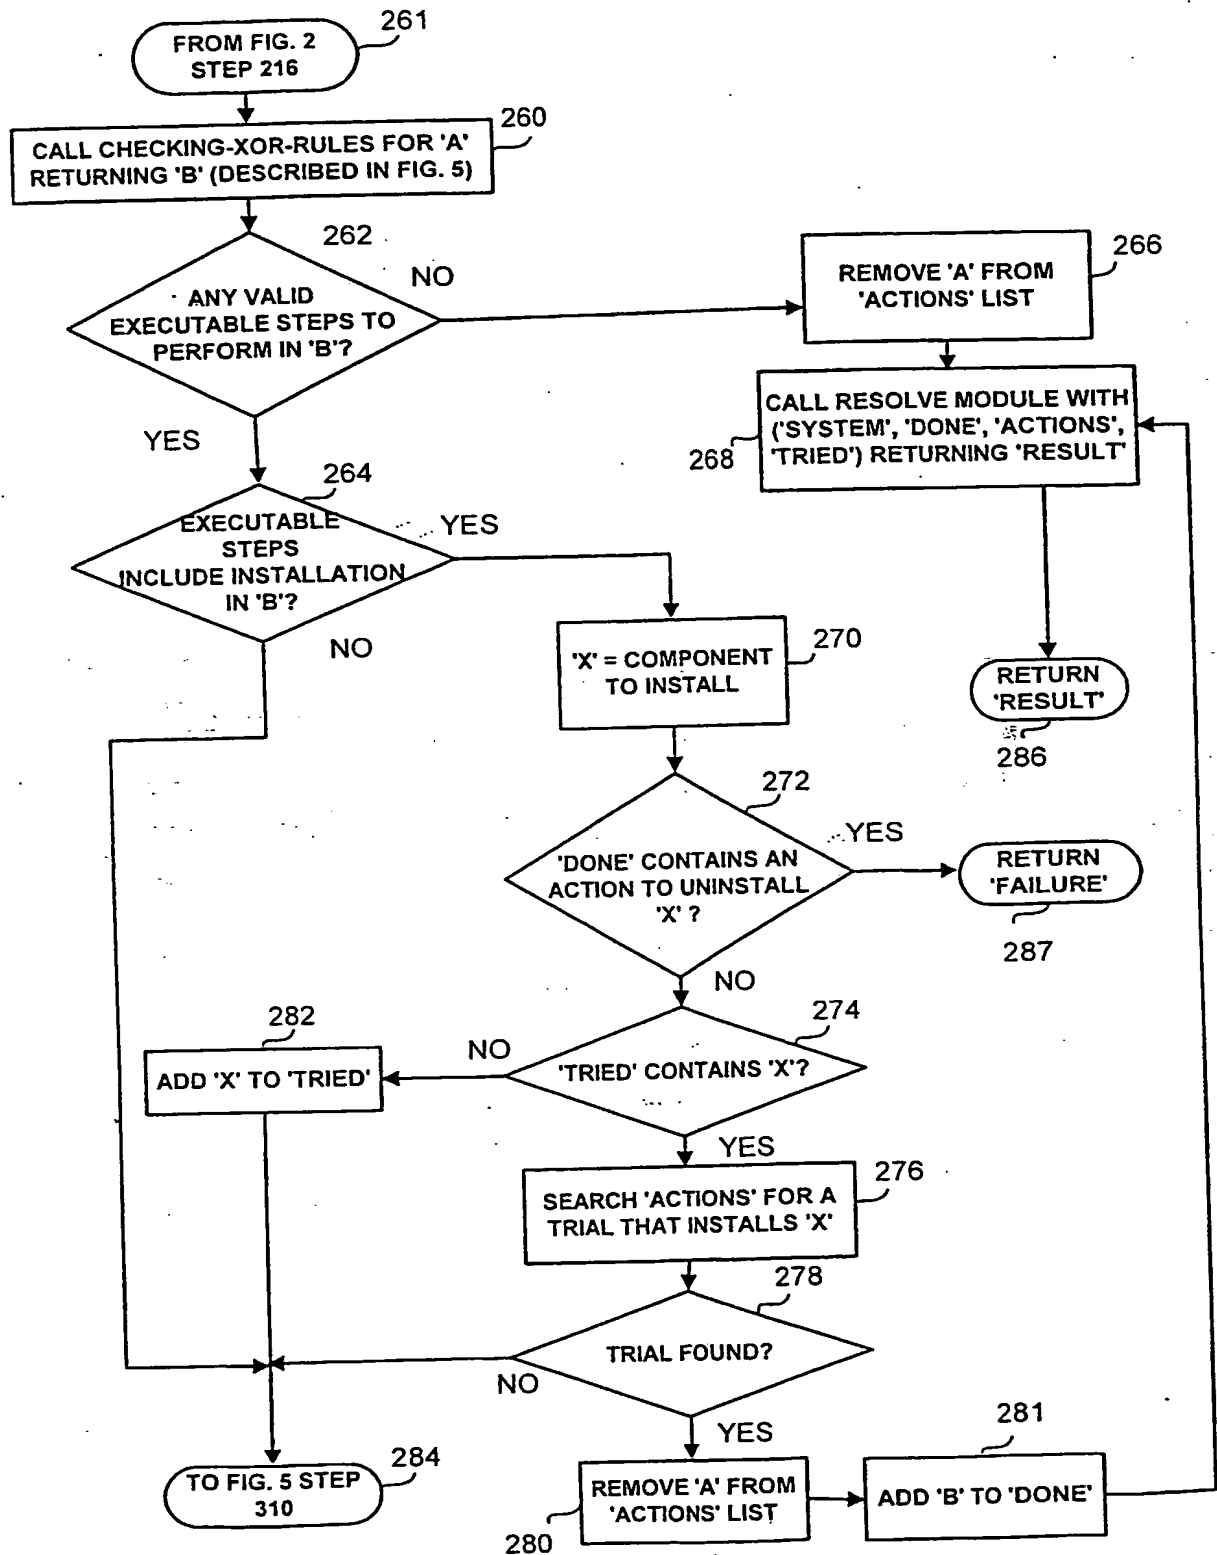


FIG. 4

5/10

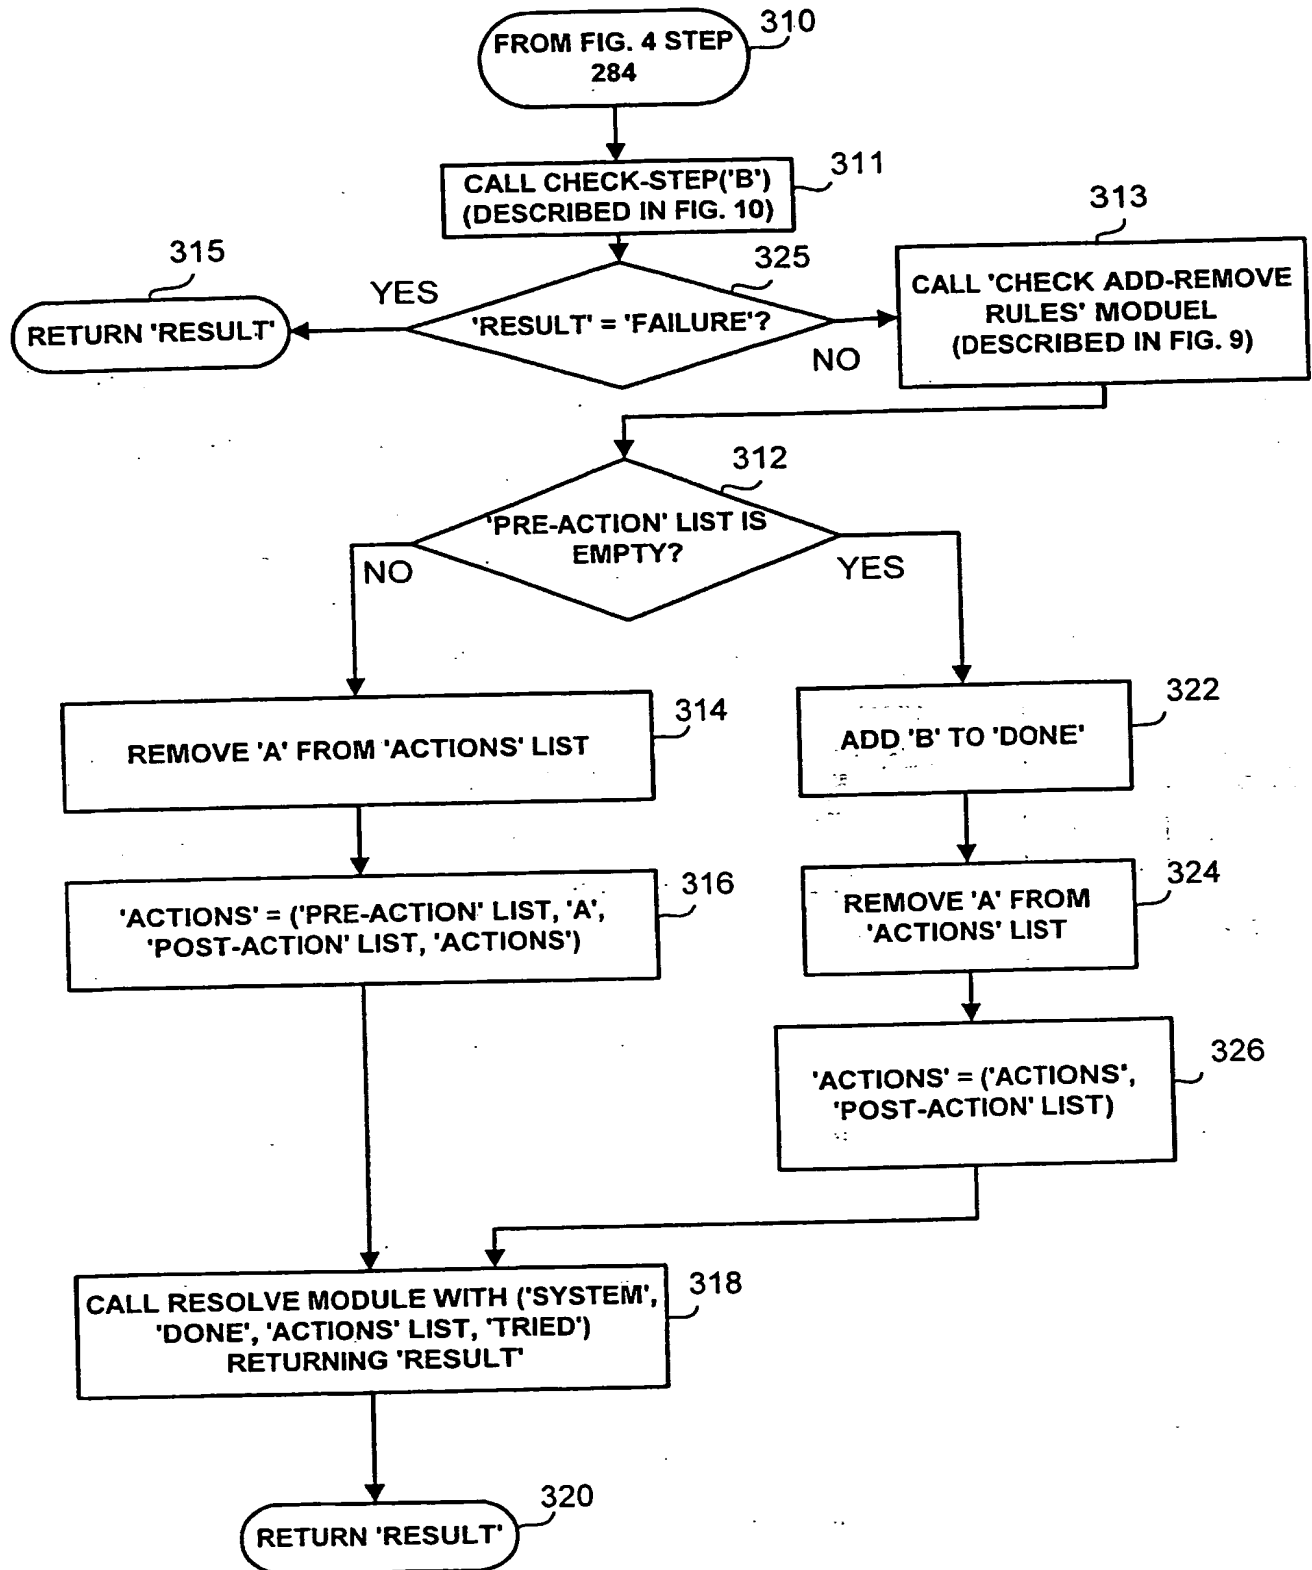


FIG. 5

6/10




FIG. 6

7/10

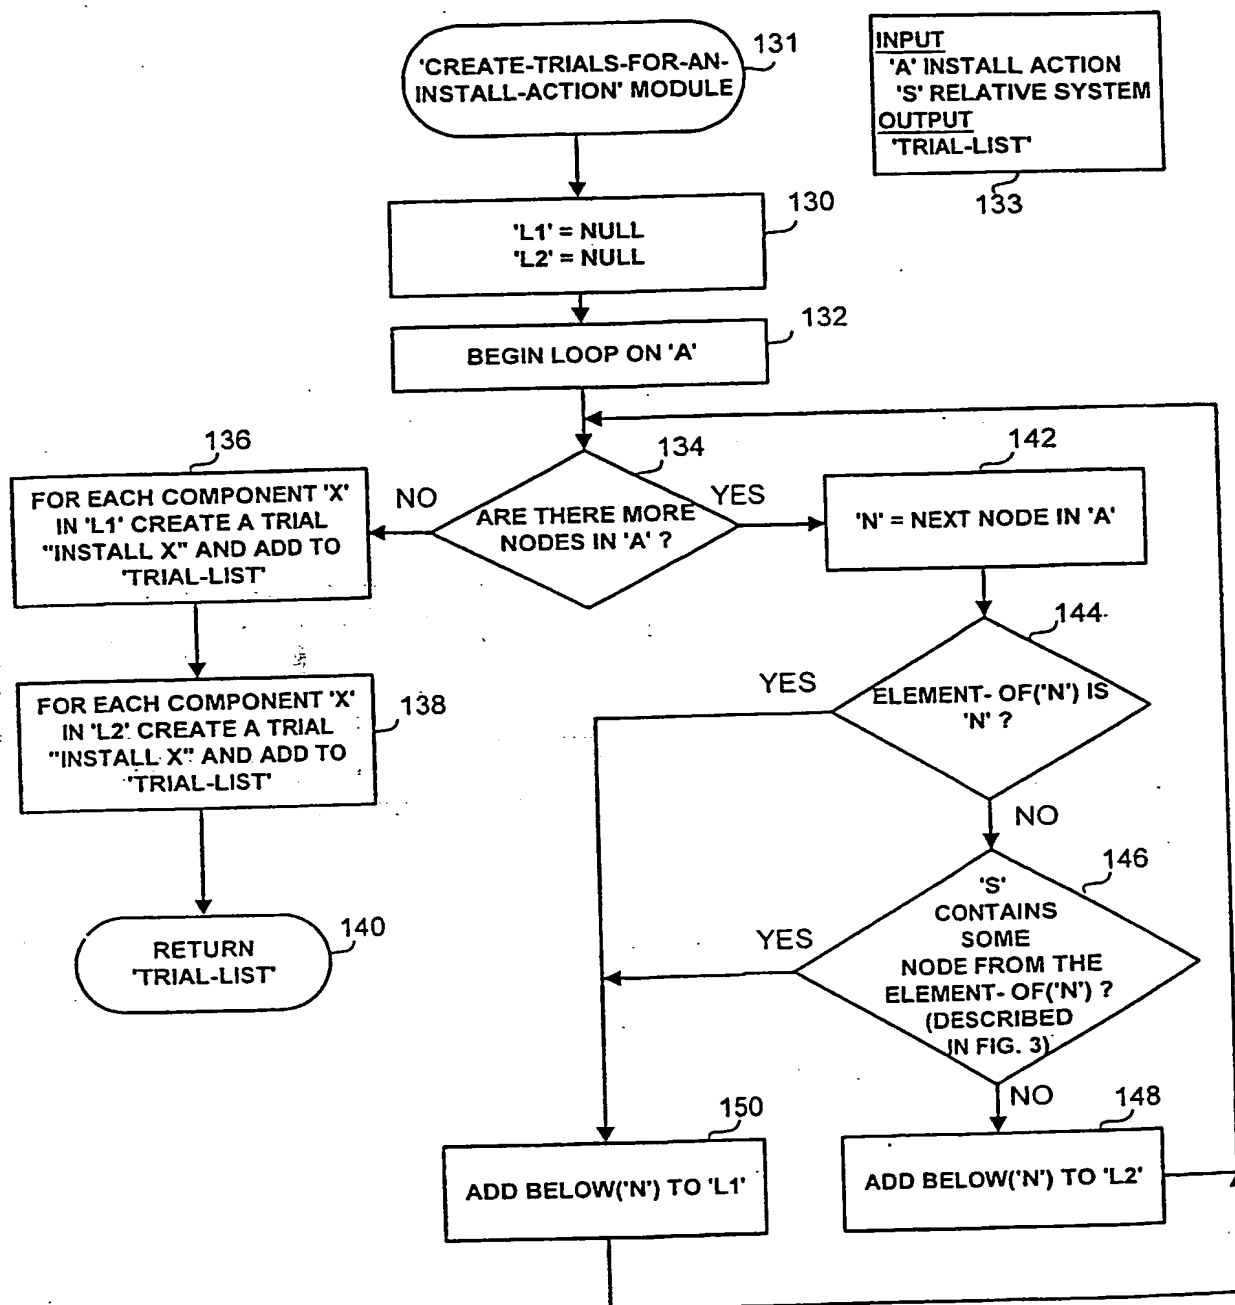


FIG. 7

8/10

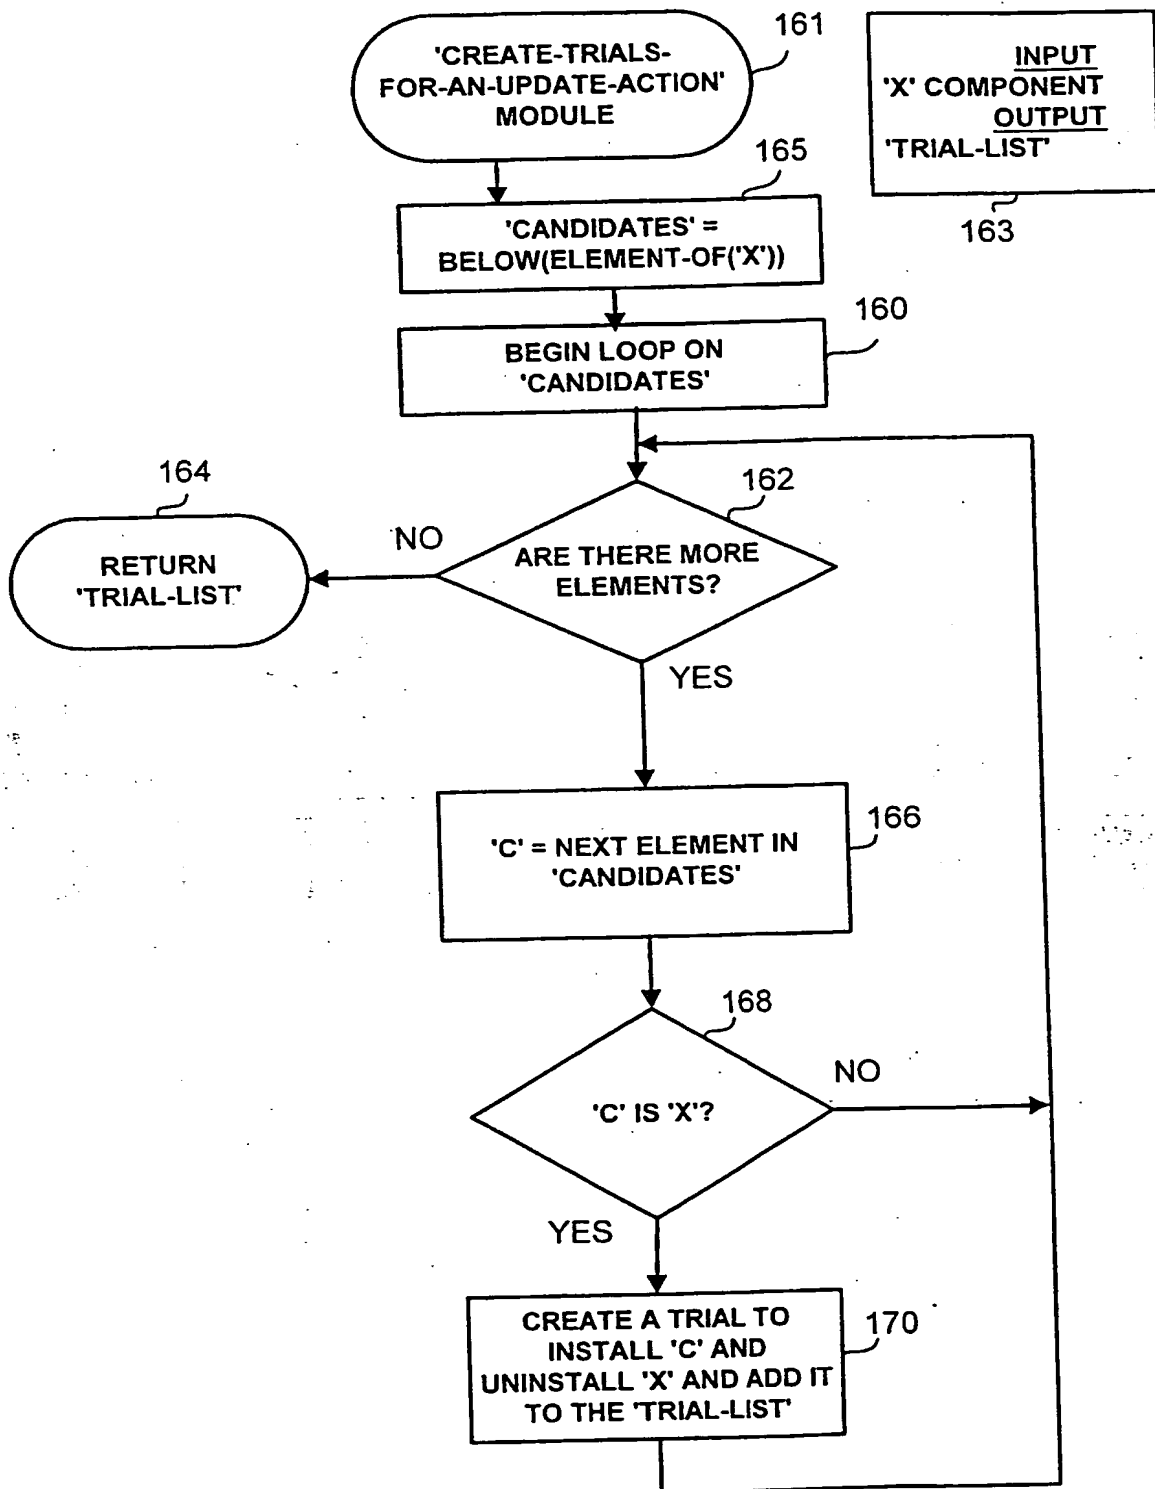


FIG. 8

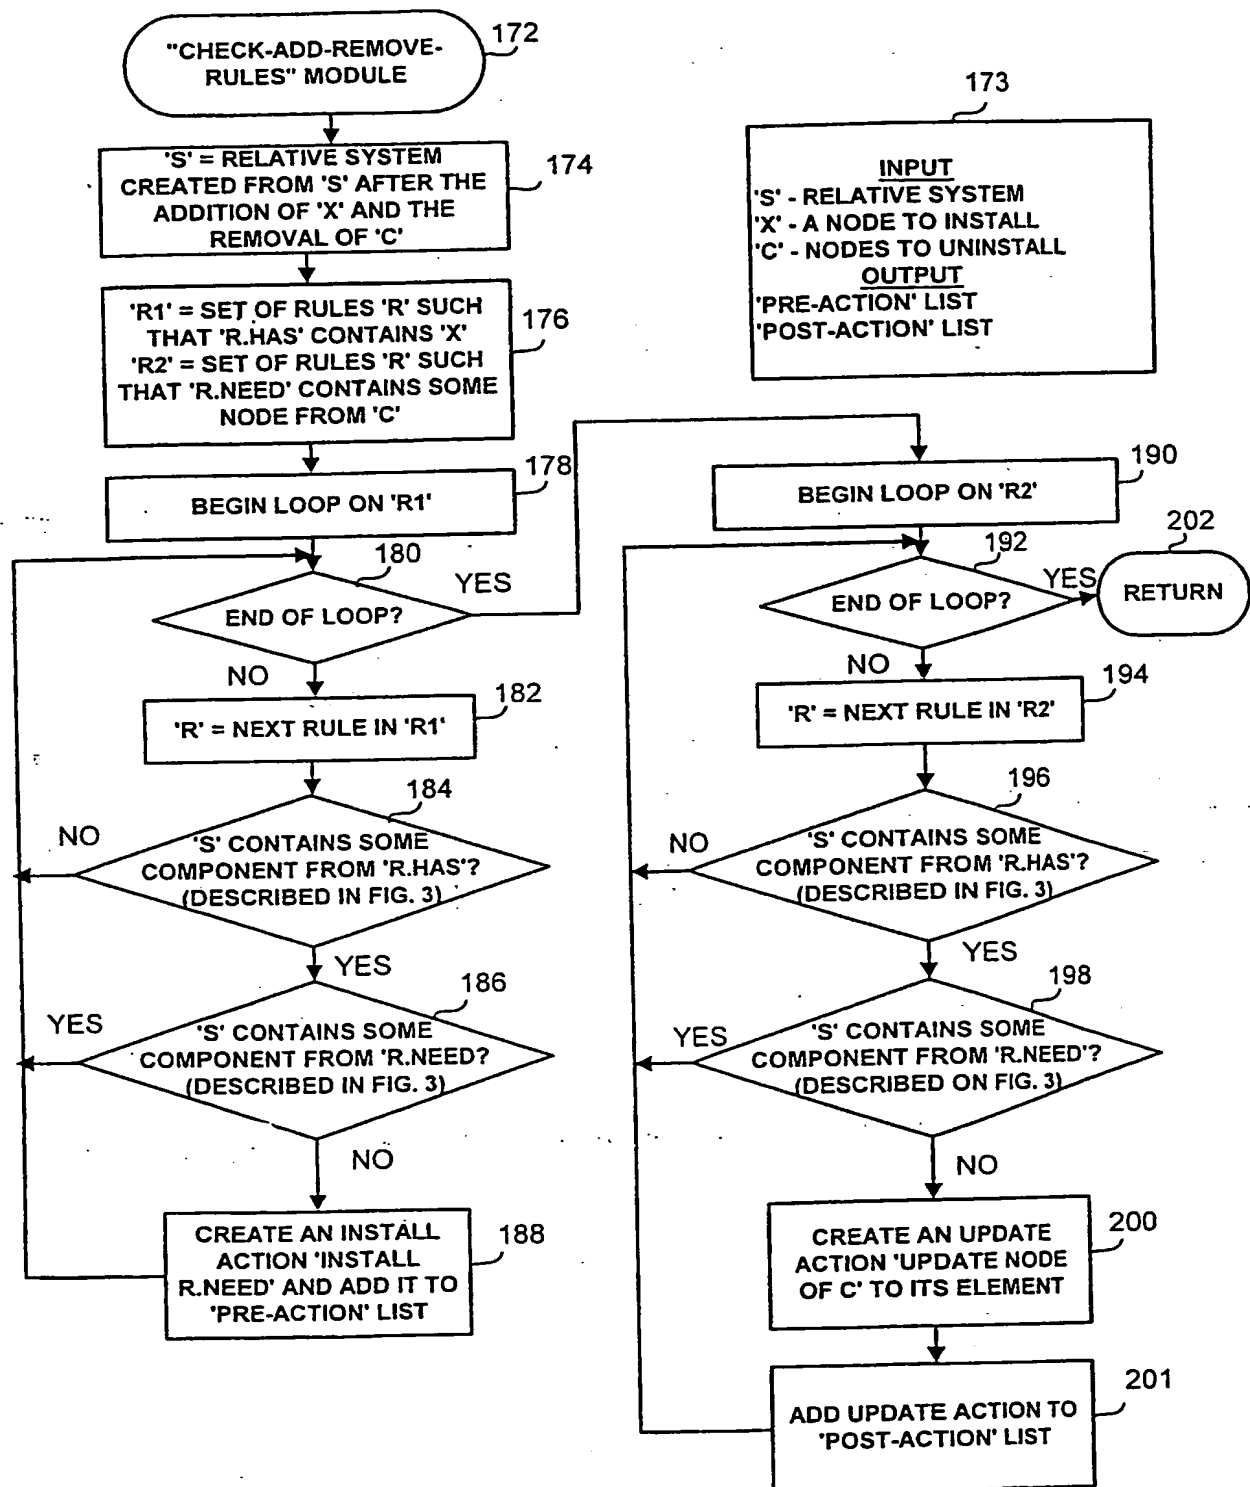
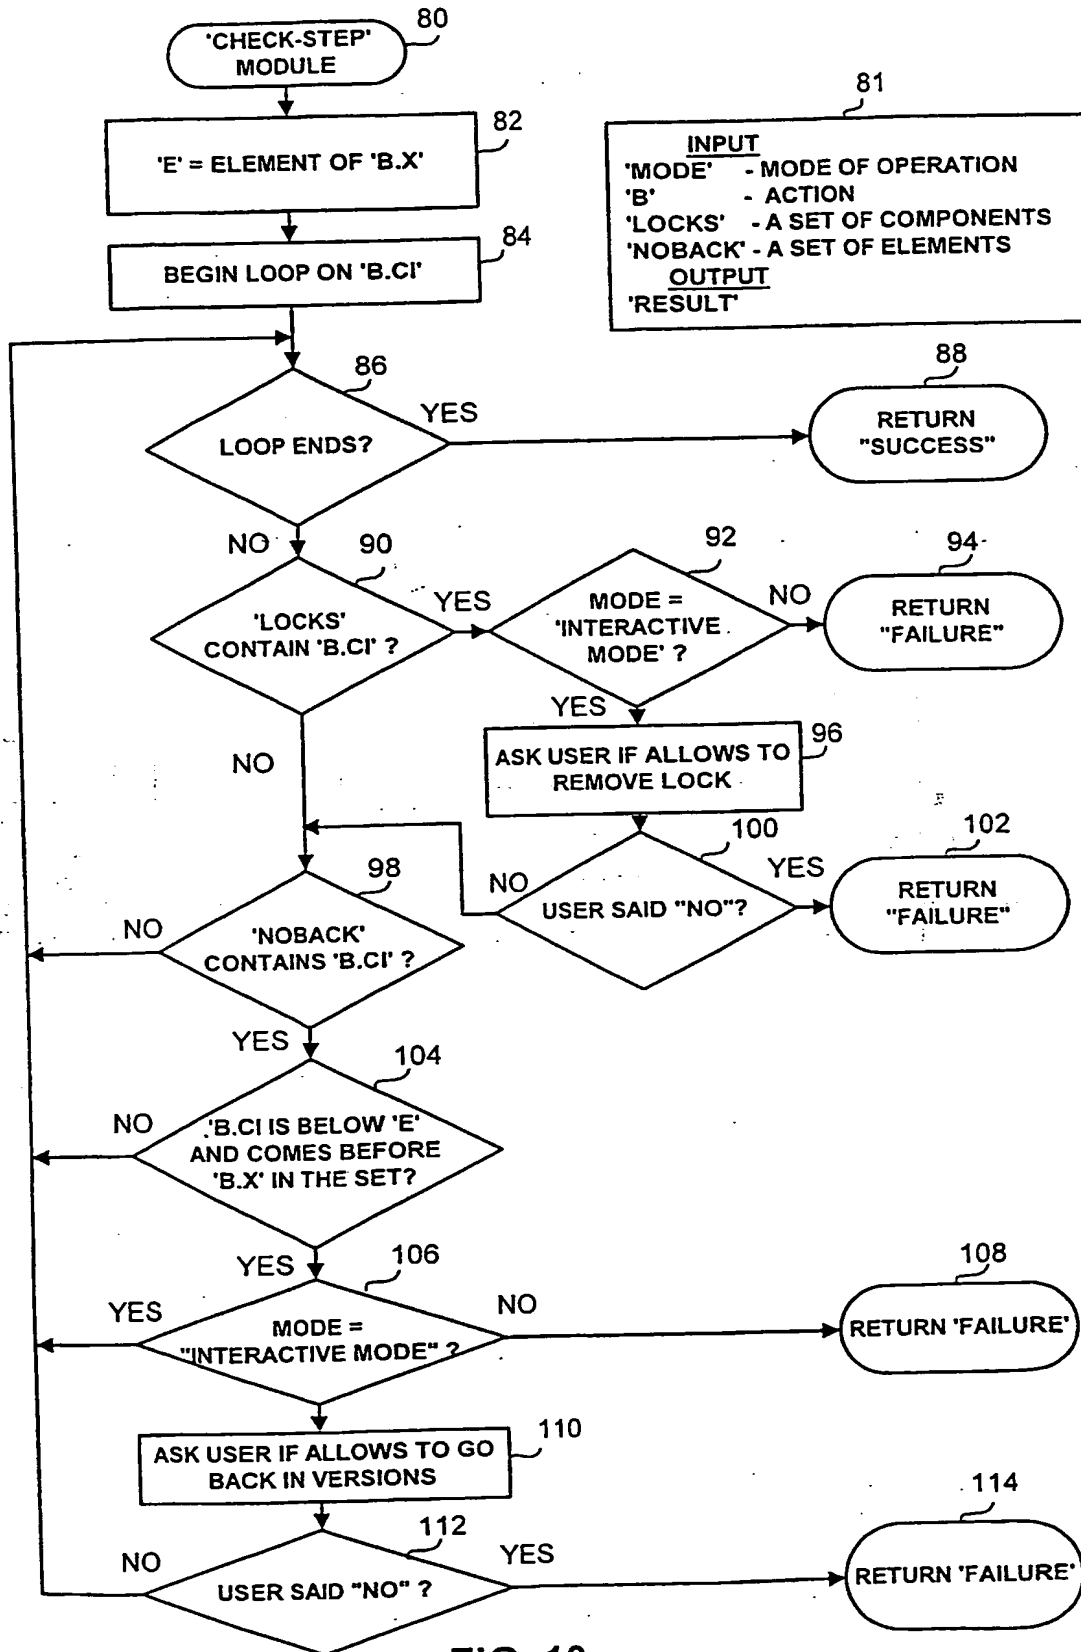


FIG. 9

10/10



(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 June 2002 (13.06.2002)

PCT

(10) International Publication Number
WO 02/046914 A3

(51) International Patent Classification⁷: G06F 9/445, 9/44

(21) International Application Number: PCT/IL01/01133

(22) International Filing Date: 6 December 2001 (06.12.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/251,918 7 December 2000 (07.12.2000) US
PCT/IL01/00208 5 March 2001 (05.03.2001) IL

(71) Applicant (for all designated States except US): ADUVA
INC. [US/US]; 19731 La Mar Drive, Cupertino, CA 95014
(US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): TE'ENI, Moddi

[IL/IL]; 3 Wiesel Street, 64241 Tel Aviv (IL). SHUFER,
Ilan [IL/IL]; 12 Hachita Street, 50200 Bet Dagan (IL).

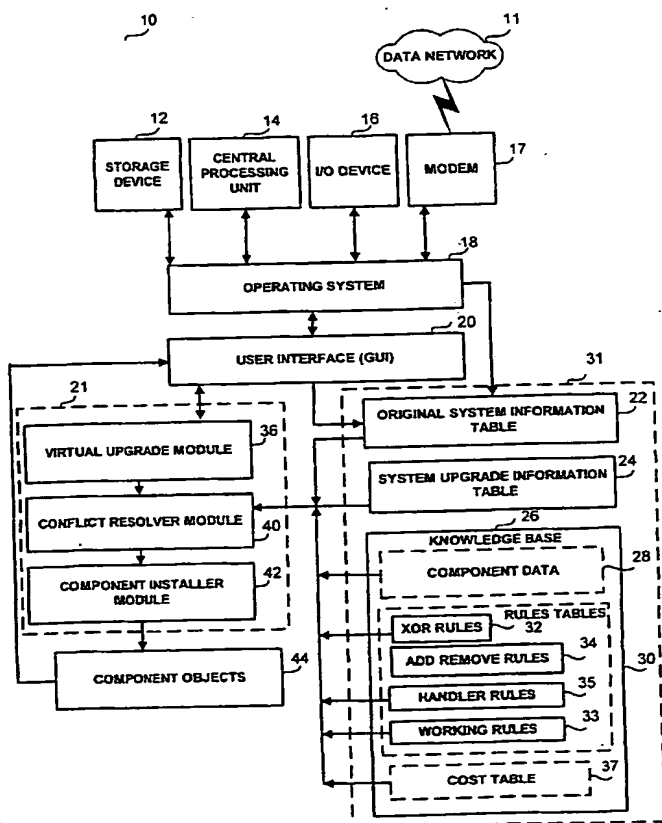
(74) Agents: AGMON, Jonathan et al.; Soroker - Agmon, 12th
floor, Levinstein Tower, 23 Petach Tikva Road, 66184 Tel
Aviv (IL).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN,
YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: IMPROVED METHOD FOR RESOLVING DEPENDENCY CONFLICTS AMONG MULTIPLE OPERATIVE ENTITIES WITHIN A COMPUTING ENVIRONMENT



(57) Abstract: An improved method to set up automatically a computer system configuration in association with a system upgrade that involves the installation of a set of operative elements necessary for the reliable operation of the same. The correct, operative combination of hardware, system and application software components is accomplished by inter-component dependency checking and improved conflict resolving procedures in association with specific dependency rules tables and an ordered database that stores definitions of components and relationships thereof.

WO 02/046914 A3



European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(88) Date of publication of the international search report:
15 August 2002

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

International Application No

PCT/IL 01/01133

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/445 G06F9/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 867 714 A (TERRELL MICHAEL R ET AL) 2 February 1999 (1999-02-02) column 3, line 15-19 column 3, line 23-25 column 12, line 1-20 column 4, line 9-19 column 14, line 15-36 column 4, line 30-42	1-3, 7-29
A	EP 0 632 371 A (XEROX CORP) 4 January 1995 (1995-01-04) page 2, paragraph 7 -page 3, paragraph 10 page 4, paragraph 29 page 3, line 18 - line 53 --- -/--	1-29

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *g* document member of the same patent family

Date of the actual completion of the international search

5 June 2002

Date of mailing of the international search report

13/06/2002

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Skomorowski, M

INTERNATIONAL SEARCH REPORT

International Application No

PCT/IL 01/01133

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	"SAFE MECHANISM FOR INSTALLING OPERATING SYSTEM UPDATES WITH APPLICATIONS" IBM TECHNICAL DISCLOSURE BULLETIN, IBM CORP. NEW YORK, US, vol. 41, no. 1, 1998, pages 557-559, XP000772211 ISSN: 0018-8689 the whole document	1-29

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/IL 01/01133

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5867714	A	02-02-1999	NONE	
EP 0632371	A	04-01-1995	BR 9402027 A	13-12-1994
			DE 69428400 D1	31-10-2001
			DE 69428400 T2	14-02-2002
			EP 0632371 A1	04-01-1995
			JP 7006026 A	10-01-1995
			US 5499357 A	12-03-1996